

Linguistic Reflexes of Well-Being and Happiness in Echo

Jiaqi Wu, Marilyn Walker, Pranav Anand and Steve Whittaker

University of California Santa Cruz

{jwu64, mawalker, panand, swhittak}@ucsc.edu

Abstract

Different theories posit different sources for feelings of well-being and happiness. Appraisal theory grounds our emotional responses in our goals and desires and their fulfillment, or lack of fulfillment. Self-Determination theory posits that the basis for well-being rests on our assessments of our competence, autonomy and social connection. And surveys that measure happiness empirically note that people require their basic needs to be met for food and shelter, but beyond that tend to be happiest when socializing, eating or having sex. We analyze a corpus of private micro-blogs from a well-being application called ECHO, where users label each written post about daily events with a happiness score between 1 and 9. Our goal is to ground the linguistic descriptions of events that users experience in theories of well-being and happiness, and then examine the extent to which different theoretical accounts can explain the variance in the happiness scores. We show that recurrent event types, such as OBLIGATION and IN-COMPETENCE, which affect people’s feelings of well-being are not captured in current lexical or semantic resources.

1 Introduction

There has recently been huge interest in well-being, with a recent review arguing that psychological well-being plays a causal role in promoting job success, physical health, and long-term relationships (Lyubomirsky et al., 2005; Kahneman, 1999). In this paper we analyze a corpus of private micro-blogs from a well-being application called ECHO, with the aim to detect, understand, and fur-

RECORDING (*Negative*): I have to clean the kitchen since it’s my chore this week, but I really don’t want to do it!
REFLECTION (*Positive*): I’m glad I did it!! The kitchen was clean and I watched the Kardashians while doing it!

RECORDING (*Positive*): I am having a lovely lunch with my two friends. We are eating at Pacific Thai. Tom yummm!!
REFLECTION (*Negative*): I miss hanging out with friends, I’ve been so busy lately.

Figure 1: RECORDING and REFLECTION of Echo

ther advance systems that can improve both short and longer-term issues with well-being.

ECHO initiates user-written reactions to daily events, called RECORDINGS, as well as subsequent REFLECTIONS on those events at points in the future (Isaacs et al., 2013).¹ Each reaction is labelled *at the time of recording or reflection* by the user, the first-person experiencer, with a *happiness* rating from 1 and 9. Note that all users’ posts and ratings are private, distinguishing this corpus from public sources like LiveJournal, where the content of posts might be influenced by considerations of self-presentation. Figure 1 shows a RECORDING and REFLECTION from two users, after binning the happiness ratings into positive and negative.

Our goal is to ground the linguistic descriptions of events that users experience, such as those in Figure 1, in theories of well-being and happiness. Without such a grounding, it is difficult for the ECHO system to make recommendations to users to improve their well-being, or to explain the relationships between different event types and well-being, or to develop a policy that can do a good job of selecting events for targeted reflection (Konrad et al., 2015; Isaacs et al., 2013). That is, for ECHO’s purposes, we need techniques that

¹The ECHO corpus is not publicly available because of the ethical agreement with ECHO users. To protect users’ privacy, the uploaded images are not stored for analysis.

not only reliably categorize a user’s scalar happiness level, but are explanatory with respect to the sources of that happiness level.

There are two principal challenges to this goal. First, different theories posit different sources for feelings of well-being and happiness. Second, the relevant computational resources for sentiment or mood are primarily lexically based, while many of the events can only be characterized well via their compositional semantics (Reschke and Anand, 2011).

Other research also shares our motivation of understanding the relationship between what people say and their levels of happiness and related moods. Mishne (2005) used a corpus of 340,000 posts from Livejournal that were self-annotated with the 40 most common moods. Lexical features alone improved classification accuracy by 6 to 15% over a balanced baseline. These results were then improved considerably (Keshtkar and Inkpen, 2009). Mihalcea and Liu (2006) experimented with the subset of *happy/sad* posts, and used conditional probability to explore the “happiness factor” of various terms, and the relationship of these terms to well-being categories such as human-centeredness and socialness. Schwartz et al. (2016) extract 5,100 public status updates on Facebook and have Turkers annotate them using Seligman’s dimensions for well-being: Positive Emotions, Engagement, Relationships, Meaning, and Accomplish (Seligman et al., 2006; Forgeard et al., 2011). They then predict each dimension with lexical and LDA topic features.

A related line of work builds lexico-semantic resources for sentiment analysis with a focus on how the participants of an event are affected by it. Goyal and Riloff (2013) bootstrap a set of patient-polarity verbs from narratives and Ding and Riloff (2016) extract event-triples from blogs that reliably indicate positive or negative affect on one of the event participants. Reed et al. (2017) take a similar approach. Deng et al. (2013) annotate how participants of an event are affected, and Deng & Wiebe (2014) show that this assists inference about the author’s sentiment towards entities or events. Balahur et al. (2012) use the narratives produced by the ISEAR questionnaire (Scherer et al., 1986) for first-person examples of particular emotions (“I felt angry when X and then Y happened”) and extract sequences of subject-verb-object triples, which they then annotate for seven

basic emotions. Choi & Wiebe (2014) use WordNet to try to learn similar patterns, and Rupenhofer & Brandes (2015) annotate synsets in GermaNet based on an event decomposition framework. Russo et al. (2015) proposed a shared task for recognition of a set of pleasant and unpleasant events from a clinical framework for well-being (MacPhillamy and Lewinsohn, 1982). Work on AFINN, SentiWordNet and the Connotation Lexicon also aim to refine existing sentiment resources to capture more subtle notions of sentiment (Feng et al., 2013; Kang et al., 2014; Baccianella et al., 2010; Nielsen, 2011).

Here we report an exploratory study where we synthesize theoretical constructs associated with well-being and happiness from different sources. We then develop several methods for characterizing events in terms of these theories. We examine the extent to which different theoretical accounts can explain the variance in the happiness scores in ECHO. We show that each theory explains a part of the variance, but that our event characterizations need to be more fine-grained. We show that several recurrent event types which affect people’s feelings of well-being, such as OBLIGATION and INCOMPETENCE, are not captured in current lexical or semantic resources.

2 Background and Motivation

ECHO is designed to encourage users to react to daily events as well as to periodically reflect on past events (Isaacs et al., 2013). Figure 2 depicts the user interface, showing a RECORDING from today, as well as prompts to reflect on events from the past. ECHO has been deployed with 134 users, in three different experiments on well-being (Konrad et al., 2016b,a). The total corpus consists of 10354 posts, where 7573 are RECORDINGS and 2781 are REFLECTIONS. While the corpus could be considered relatively small, these posts provide a window onto users’ private thoughts as opposed to what users are willing to make public on social media. In addition, the annotations for happiness are provided by the user, the first-person experimenter, and not by a third party.

Our aim is to explain users’ emotional reactions to different categories of events mentioned in ECHO posts, linking the user reactions directly to theories of well-being as exemplified in Table 1.

Influential accounts such as Appraisal Theory (Scherer et al., 2001, 1986; Ortony et al., 1990)

Row #	Source	Subtype	Affect	Example
1	Goals	Achieved	POS	I applied to an scholarship got a large chunk of my reading done and got started cramming for next test .
2		Thwarted	NEG	Wasn't able to get back in time for my class section .
3	Eudaimonics	Autonomy	POS	Good day at work had the right support and students were listening and behaving which was awesome.
4		Lack-Autonomy	NEG	Long list of things to do before going out tonight.
5		Competence	POS	After working hard and spending so many countless hours, I finally finished my project for my psych class !
6		Incompetence	NEG	My midterm was really long and I didn't finish.
7		Connection	POS	Having a nice time with my parents watching the Opening Winter Olympic Ceremony.
8		Lack or Neg-Connection	NEG	My friend needs a bone marrow biopsy and chemo.
9	Hedonics	Savouring	POS	I love home cooking! Especially if it's Italian.
		Savouring	NEG	The bus was rather packed and had a few people bump into me from where I was sitting.

Table 1: Examples of Theoretical Categories and Instantiations in ECHO



Figure 2: Screenshot of the Echo Interface

argue that success or failure in personal goals directly mediates affect. Rows 1 and 2 in Table 1. Such mediation arises because emotions have an important adaptive signaling function that serves to motivate future behaviors in relation to those goals. Row 1 provides a description from ECHO of successfully achieving goals. Appraisal theory posits that goal achievement promotes positive affect, which then serves to reinforce the relevant behavior. Row 2 provides an example of failing to achieve an important personal goal, which

is posited to promote negative affect, motivating people to modify current behaviors to change that negative outcome.

There are significant critiques of the adaptive goal-based account espoused in Appraisal theory. Appraisal theory focuses on short-term personal goals, but Eudaimonic psychologists instead focus on what determines long-term happiness. Eudaimonic theorists suggest that certain fundamental psychological needs have to be satisfied for people to experience sustained positive long-term emotions. Self-determination theory argues that there are 3 basic psychological needs: AUTONOMY, COMPETENCE and CONNECTION (Deci and Ryan, 2010; Ryan and Deci, 2000; Bandura, 1977). We add these to our inventory in Table 1 in Rows 3 to 8. According to self-determination theory, satisfaction of these basic needs results in positive emotions. Row 3 describes a good day at work. Row 5 describes feeling competent because hard work led to an achievement, and Row 7 describes feeling connected with family. On the other hand, if these basic needs are not satisfied, then negative emotions will regularly arise. For example, obligations to do things one does not feel like doing (Row 4), or a job that does not engage personal decision making or involvement (lack of autonomy) can make one feel unhappy. Similarly, people may feel unhappy due to an experience where the demands of the situation outstrip one's basic abilities, such as doing poorly on a test (lack of competence), as in Row 6. In addition, bad things happening to friends (Row 8) as well as separation from family or friends often reduces happiness (lack of connection).

In addition, there is strong evidence from SAVOURING theory (Jose et al., 2012; Bryant et al., 2011) arguing that people often experience highly positive or negative emotions arising from situations that aren’t directly goal-related, and that relate more directly to basic drives (Maslow, 1943; Elson, 2012). For example, experiences such as eating, experiencing nature, sex and physical exercise tend to engender positive emotions, whereas pain, discomfort and inactivity have the opposite effects, and these are documented in results from happiness surveys (Kahneman et al., 2004; Seligman et al., 2006). Thus while experiences such as eating may serve the survival goal of preventing starvation, avoiding starvation is unlikely to be a direct personal goal every time we eat, suggesting that such experiences are not explained by Appraisal theory. Similar arguments have been made by Lewinsohn and colleagues who have shown that encouraging people to engage in certain simple activities (shopping, mowing the lawn, driving, personal hygiene) have quite predictable effects on mood without engaging significant personal goals (MacPhillamy and Lewinsohn, 1982; Lewinsohn et al., 1985; Lewinsohn and Amenson, 1978).

3 Empirical Approach

Dataset	Pos	Neg	Total
Train	4743	3180	7923
Test	810	515	1325

Table 2: Number of Sentences for Train and Test

We start with the 10354 posts from the ECHO corpus and map happiness scores between [1, 4] to negative, and scores between [6, 9] to positive. For posts labelled 5 by the experiencer, we categorize it as negative if its REFLECTION score decreases to lower than 5, and positive if its REFLECTION score increases. We label the rest of the 5s as neutral, and leave them aside. We then have 5997 positive posts and 3573 negative posts. We randomly sample 2868 posts as training data, and 478 as test data. We keep the rest of the 6224 posts untouched for future work. Then we split the posts into sentences. Table 2 shows the splits for each class.

We first test the separability of the positive and negative sentences with an SVM classifier from Weka 3.8, using as baselines only unigrams and LIWC (Pennebaker et al., 2001) as features. Results for these baseline classifiers are in Table 3,

Features	Metric	Pos	Neg	All
UniGram	Prec	0.75	0.66	0.72
	Rec	0.81	0.59	0.72
	F1	0.78	0.62	0.72
LIWC	Prec	0.72	0.72	0.72
	Rec	0.89	0.45	0.72
	F1	0.80	0.55	0.70

Table 3: Weighted Metrics for SVM on Test

Unigram	LIWC
fun	affect,posemo,leisure
good	affect,posemo,drives,reward
we	we,social,drives,affiliation
lunch	bio,ingest
glad	affect,posemo
want	cogproc,discrep
why	interrog,cogproc,cause
need	cogproc,discrep
no	negate
not	negate,cogproc,differ

Table 4: The most informative UniGram features weighted by Information Gain

LIWC	Words
negemo	stress*, sad, sick, hate
posemo	fun, well, great, love
negate	dont, didnt, no, cant, havent
anger	hate, frustrat*, annoying
i	i, my, me, im, myself
differ	but, not, really, didnt
leisure	fun, game*, relax*, family
discrep	want, need, would, should
sad	sad, miss, hurt*, missed
risk	stop, problem*, avoid*
anx	stress*, nervous, worried
ingest	food*, dinner*, lunch*
body	sleep, slept, stomach*
insight	feel, know, think, found
affiliat	we, friends, friend, love
reward	good, got, get, great
feel	feel, feeling, felt, hard
family	family, mom, sister*, dad
we	we, our, us, weve, lets

Table 5: The most informative LIWC features ranked by Information Gain.

illustrating that the positive and negative classes can be separated with F1 above .70, and that both unigrams and LIWC perform worse on the negative class.

However, as discussed above, the word level representations of the features in the baselines do not help us with our goal to understand how linguistic descriptions of events that affect well-being map onto theoretical constructs. Table 4 and Table 5 provide the most informative UniGrams

Well-Being	Frames	Example Lexical Units
Goal	Desiring, Intentionally_Act, Purpose	want, feel like, hope, wanted, wish, do, did, done, doing, does, plan, purpose, in order, intention, goals
Autonomy & Obligation	Being_obligated, Required_event, Avoiding, Inhibit_movement, Have_as_requirement, Complaining	complain, grumble, complaints, have to, had to, should, having to, need, get to, had to, have to, got to, should, avoid, ducking, take, need, needed, requires
Competence	Activity_done_state, Attempt, Capability, Bungling, Difficulty, Practice, Activity_finish, Accomplishment	finished, trying, try, tried, effort, attempt, efforts, can, could, exercise, practice, rehearsal, exercising, able, ability, unable, messed up; ruined; screwed up, ruin, hard, difficult, easy, tough, easier challenging, impossible, a breeze, hardest, finish, finishing, completed, accomplished, achieve
Connection & Lack-of Connection	Death, Forming_relationships, Social_event, Kinship, People, People_by_residence, Telling, Communication_response	birthday, married, divorce, befriend, dinner, social, party, picnic, mom, family, parents, sister, cousin, told, tell, informed, people, girl, man, roommate, reply, answers, answer, reacted
Savouring	Emotions_of_mental_activity, Feeling, Annoyance, Desirability, Food, Chemical-sense_description, Ambient_temperature, Emotions-by-stimulus, Stimulus_focus, Intoxicants, Communication_noise, Experiencer_Focus, Perception_experience, Biological_urge, Death	enjoyed, like, hate, glad, annoyed, cry, yelled, whooped, honked, irritated, feel, feeling, yummy, alcohol, weed, drugs, dope, see, felt, seeing, hear, experience, senses, experiences, taste, feel, delicious, tasty, sweet, food, coffee, bread, cheese, good, bad, great, better, best, horrible, worst wonderful, weird, nice, relaxing, annoying, interesting, sad, weird enjoyable, comforting, entertaining, unpleasant, hilarious, rest, relaxation, exhilarating, tiring, nicer, disturbing, disappointing, embarrassing, irritating, upsetting, heartbreaking, consoling, tedious, traumatic, chilling, calming, frightening touching, pleasure, satisfying, fascinating, tired, exhausted, sleepy, hungry, nauseated, horny

Table 6: Frame Categories and Associated Well-Being Classes.

and LIWC categories. We cannot recommend to an ECHO user that they should for example, try to use the word *why* less (Row 7) because it is correlated with negative feelings, or try to use less negation (Rows 9 and 10). It is difficult to associate these features with well-being classes. Even in cases where the words seem to be strongly related to a well-being category, a single word typically fails to provide enough information, e.g., “it was **fun** talking to him” and “worked on a **fun** project” belong to different well-being classes. Moreover, the mapping of LIWC categories to words are many-to-many, e.g. the “discrep” category contains words related to both Goals and Autonomy. We posit that we need compositional semantic features to ground our a Well-Being classification of events.

We thus explore two different methods for mapping these well-being event categories into lexical descriptions, one of which is top-down and the other which is bottom-up. Our top-down method is based on mapping general event types from FrameNet to the theoretical categories enumerated in Table 1. We take frame specific features for each theoretical category from the lexical units for each frame. For example, GOALS are often dis-

cussed in terms of specific frames from the Desiring and the Intentionally_act classes, as shown in the first two rows of Table 6.

We show that FrameNet features do provide an interesting level of generalization but much of the compositional semantics of events is still missing from this characterization (Section 4). Thus, our bottom-up method applies the AutoSlog linguistic-pattern learner to induce lexically-grounded predicate patterns from the ECHO data (Section 5). We show how many light verbs acquire a specific semantics with their arguments, and how common events like “Talking” are separated into positive and negative events depending on whether they are “Talking about” or “Talking with”.

4 Frames and Well-Being

Table 6 provides our posited mapping from frame categories to the appraisal category of GOALS as well as to the eudaimonic categories of AUTONOMY, COMPETENCE and CONNECTION, and to the hedonic category of SAVOURING. To develop features related to these frame categories, we apply SEMAFOR (Das et al., 2013) to label the ECHO posts with their corresponding frames using

FrameNet 1.5 (Baker et al., 2015; Baker, 2014). We partition frame features into subsets corresponding to the different theoretical constructs as defined in Table 6. We acknowledge that our mapping may not be perfect, and that some frames could conceivably be categorized as both goal related and eudaimonic.

Features	Metric	Pos	Neg	All
GOALS	Prec	0.62	0.49	0.57
	Rec	0.94	0.09	0.61
	F1	0.75	0.15	0.51
EUDAIMONIC	Prec	0.63	0.58	0.61
	Rec	0.93	0.16	0.63
	F1	0.75	0.25	0.56
SAVOURING	Prec	0.61	0.44	0.55
	Rec	0.97	0.04	0.61
	F1	0.75	0.08	0.49
ALL FRAMES	Prec	0.69	0.74	0.71
	Rec	0.91	0.38	0.70
	F1	0.78	0.50	0.67

Table 7: Coverage of Different Theoretical Categories.

We train an SVM with each feature subset, and evaluate the models on our test set, with results in Table 7. The general ALL FRAME feature is also listed for comparison. The .67 F1 of FRAME is slightly lower than LIWC in Table 3, but in our view, more interpretable. In addition, the average count of FRAME features per sentence is an order of magnitude less than LIWC features (hence, much less than unigram features), suggesting the targeted power of these features. See Table 8. We posit that FRAMES are thus more discriminative than LIWC for well-being classes, and that FRAME features are more naturally categorized into well-being categories at a semantic level.

Features	Dataset	Pos	Neg	Total
UniGram	Train	8.5	9.9	9.1
	Test	8.1	9.8	8.7
LIWC	Train	25.4	31.4	27.8
	Test	23.8	30.6	26.4
ALL FRAMES	Train	2.7	5.2	3.7
	Test	3.3	4.0	3.6

Table 8: Average Feature Counts for Sentence

The Goals section of Table 7 shows that Appraisal theory does well at predicting positive events, but performs poorly for negative events, primarily due to low recall. All features achieve

Features	Metric	Pos	Neg	All
AUTONOMY	Prec	0.0	0.39	0.15
	Rec	0.0	1.0	0.39
	F1	0.0	.56	0.22
COMPETENCE	Prec	0.56	0.58	0.60
	Rec	0.98	0.04	0.61
	F1	0.76	0.07	0.49
CONNECTION	Prec	0.62	0.58	0.60
	Rec	0.97	0.06	0.62
	F1	0.76	0.11	0.49

Table 9: Results for Individual Eudaimonic Categories.

good F1 for the positive class, but not the negative class. This is consistent with the results in Table 3.

The EUDAIMONIC features include Autonomy & Obligation, Competence and Connection. The SVM trained with just eudaimonic features produces the highest F1 score for the negative class, highlighting the role of eudaimonic related events in negative well-being. See Table 7. The results for an breaking eudaimonic into its constituent categories is in Table 9. The results show that most of our autonomy categories are related to negative autonomy, to obligations that cause feelings of negative well-being. On the other hand, the results indicate that competence and connection play a large role in positive well-being.

The top 25 most informative frame features are illustrated in Table 10 (out of 639 instantiated in ECHO). These illustrate general events for well-being, but compositional differences, such as “spending my nights by the side of my textbook” and “spending my nights with friends” are not captured. The first “spend (time)” evokes the theoretical construct of obligation, while “spend (time with)” is related to connection.

5 Linguistic Pattern Learning

We also apply Autoslog-TS, a weakly supervised linguistic-pattern learner as a way of learning some compositional patterns. Autoslog only requires training documents labeled broadly into our two classes of POSITIVE or NEGATIVE. The learner uses a set of syntactic templates to define different types of linguistic expressions. In general, this method tends to produce high precision (and potentially low recall) markers of the particular classes that can seed further hypothesizing.

The left-hand side of Table 11 lists example pattern template and the right-hand side illustrates

Well-Being	Frame	Affect	Example
GOALS	Desiring Intentionally_act	POS NEG	I think it went well and I hope I did a good job. My midterm was really long and I didn't finish.
AUTONOMY	Being_obligated Required_event	NEG NEG	I'm mad that I had to drive all the way to Fresno. I need to stay awake and listen, but it 's hard.
COMPETENCE	Capability Attempt	POS NEG	I feel so empowering whenever I'm able to help others. Tried to chat with some people online, did n't work out.
CONNECTION	Kinship Telling	POS NEG	My mom and I hung out and walked around for 6 hours . I wonder how much they will tell me my teeth are bad today.
SAVOURING	Chemical-sense_description Food	POS POS	Yummy burgers and sides. Made homemade ice cream with my husband:... cookie dough

Table 10: Top Frame Categories and Associated Well-Being Classes.

a specific lexico-syntactic pattern (**in bold**) that represents an instantiation of each general pattern template for learning well-being patterns in our data.²

In order to enable selection of particular patterns, AutoSlog-TS computes statistics on the strength of association of each pattern with each class, i.e. $P(\text{POSITIVE} | p)$ and $P(\text{NEGATIVE} | p)$, along with the pattern's overall frequency. We define two tuning parameters for each class: θ_f , the frequency with which a pattern occurs, θ_p , the probability with which a pattern is associated with the given class. AutoSlog lets us systematically explore tradeoffs with precision and recall. Here we select θ_f and θ_p to optimize F1 on our test set. For more detail, see (Riloff, 1996; Oraby et al., 2015).

Our primary interest here is Autoslog's ability to learn compositional patterns. Autoslog can, in principle, provide three kinds of information: i) it can provide supplement the lexical units for a given frame; ii) it can supplement the frames in a well-being category; and iii) it can reveal reliable markers of mood that well-being categories do not capture. Because our interest in frames is ultimately as a way of relating well-being categories with linguistic signals, we will not distinguish (i) and (ii) here.

Here we discuss all patterns with a $\theta_p > .7$. Several lexicosyntactic patterns fit within our well-being categories but are not captured by frames, while as expected there are overlaps between FrameNet and Autoslog as well. Examples are listed in Table 12. One large class includes straightforward lexical patterns: FINISHED, FIN-

²The examples are shown as general expressions for readability, but the actual patterns must match the syntactic constraints associated with the pattern template.

ISH, and FINALLY which we associate with feelings of competence. Verbal patterns with EAT and ATE indicate savouring, with NOT_EAT reliably marking negative sentences. The frames also show many specific types of food (cake), and we use a comprehensive list from DBpedia (Lehmann et al., 2014) to collapse all these to the general type FOOD, allowing us to develop patterns such as MADE_FOOD.

Autoslog also discovers many patterns syntactically linking content (nouns and verbs) and function words (e.g., prepositions and light verbs). It thus furnishes a ready source for multi-word, partially compositional expressions of positivity or negativity. In what follows, we provide some examples (note that in the patterns below, expressions in brackets are used to indicate expressions not part of the pattern that correlate with it in the data).

There are 262 positive patterns of the form Verb/Noun + "with", e.g. TALKED_WITH, DINNER_WITH, BREAKFAST_WITH, STUDYING_WITH, PLAYED_WITH, TIME_WITH, MET_WITH, SHOPPING_WITH, COFFEE_WITH, all of which describe activities that involve connection. There are also 100 negative patterns of this form, which are much more heterogeneous, involving both negative social experiences (ARGUMENT_WITH, DRAMA_WITH, INFURIATED_WITH, FIGHT_WITH), but also various problematic events (STRESSED_WITH, DIFFICULTIES_WITH, DISSATISFIED_WITH) and instruments for negative events (STOP_WITH, POISONING_WITH). Moreover, while the positive patterns cover 523 sentences in the data, the negative patterns cover only 133.

There are 62 patterns involving the string

	Pattern Template	Example Instantiations
1	<subj> PassVP	<I> am so relaxed after getting to sleep in and rest all morning.
2	<subj> ActVP	When it does happen, <I> feel energized because IT IS a special experience to me.
3	<subj> ActVP Dobj	<I> enjoy his efforts lately to make me happier.
4	<subj> ActInfVP	Found some some stuff but I AM not sure if <I> want to keep them.
5	<subj> PassInfVP	2 of <my housemates> were supposed to clean on Tuesday and they still haven't.
6	<subj> AuxVP Dobj	We ate and <We> had a glass of my favorite wine.
7	<subj> AuxVP Adj	<All of the colors> are so much more vibrant.
8	ActVP <dobj>	Cannot wait to study while eating <this> .
9	InfVP <dobj>	Just realized I forgot to turn in <my homework> .
10	ActInfVP <dobj>	I really need to start <my hw> sooner...
11	Subj AuxVP <dobj>	IT IS the Super Bowl today and THERE IS <a party> at my house.
12	NP Prep <np>	Driving in <the rain> is scary.
13	ActVP Prep <np>	Almost as if I forgot something terribly important or I messed up <something> important in my life.
14	PassVP Prep <np>	And I feel like I did but just this once I messed up and I might be punished for <it> .
15	InfVP Prep <np>	Felt amazing to be done with <finals> !
16	<possessive> NP	<Her> attitude is not working anymore.

Table 11: AutoSlog-TS Templates and Example Instantiations

“talk”, 32 positive (71 items) and 30 negative (66 items). The positive ones strongly indicate connection (e.g., TALK_WITH, HAVE_TALK, REMEMBER_[TO]-TALK, GOT_[TO]-TALK, TALK_THROUGH). In contrast, the negative index

either the obligation to talk (e.g., TRYING_TALK, NEED_TALK, HAVE_[TO]-TALK) or a failure to talk (e.g., NOT_TALK_TO, NOT_WANT_TALK, STOP_TALKING).

There are 36 patterns with the string ‘go’, 12 positive (16 items) and 24 negative (40 items). There are 34 patterns involving the past tense form “went”, which reverses the polarity to 25 positive patterns (273 items) and 9 negative (9 items). Across the two versions of the lemma, the positive patterns provide several expressions for savouring (WENT/GO_ON/FOR [a walk, a hike, a ride], WENT/GO_SHOPPING/SWIMMING, WENT/GO_TO [the mall, a movie]). For the negative, the predominance of ‘go’ comes from the fact that they are largely negated (NOT_GO_TO [the movies]) or in infinitive contexts that suggest obligation ([HAVE TO]-GO_TO [class], [HAVE TO]-GO_WORK). Similarly, the positive class contains 9 patterns with ‘bought’ and 1 with ‘buy’ (ENTICED_[TO]-BUY) and the negative class has 6 patterns with ‘bought’ and 16 with ‘buy’, all emphasizing buying necessities (BUY_GROCERIES/TICKET, NEED/WANT_BUY, NOT_BUY) Thus, even though these expressions all involve the same verbs and prepositions, the surrounding environments, as reflected in the form of the verb, split between positive and negative sentence classes.

There are 73 bigram patterns of the form NEW_X, 56 positive (83 items) and 17 negative (21 items). In general, the positive ones describe new objects – SHIRT, SHEETS, COMPUTER, CLOTHES, TEA – and acquaintances (NEW_FRIEND), thus encompassing both Connection and possibly Savouring. In contrast, the negative patterns describe changes to routines – HABITS, school QUARTER, PROFESSOR, LIVING [conditions], or SCHEDULE – which are likely to engender a sense of instability, and hence be Eudaimonically negative.

Thus, these patterns illustrate that Autoslog can serve as a high-precision method of building additional patterns – especially compositional ones – for a given well-being category.

6 Conclusions and Future Work

In this paper, we have advanced a synthetic categorization of the sources for well-being and happiness. We have used a corpus of private microblogs from the ECHO application to explore how well we can map linguistic expressions of well-being to this classification. We have shown that

Prob.	Freq.	Pattern and Text Match	Sample Post
Positive Example Patterns			
1.00	11	ActVp Prep <NP> (WENT ON)	I just went on a hike this is the best thing ever.
1.00	7	<subj> ActVP Dobj (MADE FOOD)	Made a German pancake for breakfast.
1.00	7	NP Prep <np> (CATCHING WITH)	Catching up with old friends!
1.00	7	ActVP <dobj> (USED)	Used the Laurel’s Kitchen Bread Book recipe.
1.00	6	ActVP Prep <np> (GOT OFF)	Got off work.
1.00	4	NP Prep <np> (TALK WITH)	Having a really nice talk with my aunt.
0.95	18	ActVP <dobj> (FINISHED)	Finished my paper.
0.78	39	ActVP <dobj> (TOOK)	Took a walk after class and truly enjoyed the outdoors!
0.78	25	<subj> ActVP (ATE)	We ate and had a glass of my favorite wine.
0.73	11	InfVP Prep <np> (SPEND WITH)	Happy to simply spend time with friends.
Negative Example Patterns			
1.00	9	InfVP <dobj> (AVOID)	Better buy ... in smaller packaging to avoid wasting again.
1.00	8	ActVP <dobj> (USE)	All she did was use water and wipe a few corners.
1.00	7	InfVP <dobj> (STOP)	I need to stop smoking.
1.00	6	<subj> ActVP Prep <np> (NOT TALK TO)	And now my bf is busy and can’t talk to me.
1.00	5	<subj> ActVP Dobj (TEXTED ME)	He texted me finally but then he randomly stopped.
1.00	5	<subj> ActVP (NOT SLEEP)	Have to get up early and I can’t sleep .
1.00	4	ActVP <dobj> (NOT FIND)	I did not find the time to finish my homework.
0.82	14	<subj> ActVP (REALIZED)	I JUST realized that I have to go tomorrow.
0.81	13	<subj> ActVP (TAKE)	Since I take around 35 minutes to get ready, I missed ...
0.80	20	ActVP <dobj> (TOLD)	Told my mom about my grades.

Table 12: Examples of Characteristic ECHO Patterns using AutoSlog-TS Templates

FrameNet provides useful generalizations, while the linguistic pattern learner AutoSlog illustrates the details and challenges of the compositional nature of user’s descriptions of their daily experiences. Moreover, we have demonstrated that, independently, each of these methods can produce performance similar to that of conventional lexical methods with a feature space that is smaller, and, in the case of FrameNet features, psychologically grounded. Our Autoslog exploration moreover reveals a way of exploring the space of patterns that our FrameNet mapping has missed. In future work, we aim to automatically combine these two methods and bring the Autoslog patterns under the well-being categorization we have advocated here. We also plan to investigate new models with the untouched 6224 Echo posts, as well as larger public corpus like LiveJournal.

In addition, we plan to explore the source of the fact that there are more positive patterns (both as types and the tokens they capture) than the negative ones, which directly relates to the lower Neg recall for all classifiers we tested. While we could not find any clear reason in our examination of the data, this asymmetry may indicate that markers of negativity are more syntactically distributed than

our current list of patterns looks for, or perhaps less linguistically reliable.

Acknowledgments

This research was partially supported by NSF Robust Intelligence #IIS-1302668-002 and NSF HCC #IIS-1321102.

References

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *LREC*. volume 10, pages 2200–2204.

Colin Baker, Nathan Schneider, Miriam R L Petruck, and Michael Ellsworth. 2015. [Getting the roles right: Using framenet in nlp](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorial Abstracts*. Association for Computational Linguistics, Denver, Colorado, pages 10–12. <http://www.aclweb.org/anthology/N15-4006>.

Collin Baker. 2014. [Framenet: A knowledge base for natural language processing](#). In *Proceedings of Frame Semantics in NLP: A Workshop in Honor of Chuck Fillmore (1929-2014)*. Association for Computational Linguistics, Baltimore, MD, USA, pages 1–5.

- <http://www.aclweb.org/anthology/W/W14/W14-3001>.
- Alexandra Balahur, Jesus M. Hermida, and Andres Montoyo. 2012. Building and exploiting emotinet, a knowledge base for emotion detection based on the appraisal theory model. *IEEE Trans. Affect. Comput.* 3(1):88–101.
- A. Bandura. 1977. Self-efficacy: toward a unifying theory of behavioral change. *Psychological review* 84(2):191.
- Fred B Bryant, Erica D Chadwick, and Katharina Kluwe. 2011. Understanding the processes that regulate positive emotional experience: Unsolved problems and future directions for theory and research on savoring. *International Journal of Wellbeing* 1(1).
- Yoonjung Choi and Janyce Wiebe. 2014. +/-effectwordnet: Sense-level lexicon acquisition for opinion inference. In *EMNLP*. pages 1181–1191.
- Dipanjan Das, Desai Chen, André F. T. Martins, Nathan Schneider, and Noah A. Smith. 2013. Frame-semantic parsing. *Computational Linguistics* XX(YY):WW–ZZ.
- Edward L Deci and Richard M Ryan. 2010. *Self-determination*. Wiley Online Library.
- Lingjia Deng, Yoonjung Choi, and Janyce Wiebe. 2013. Benefactive/malefactive event and writer attitude annotation. In *ACL*. pages 120–125.
- Lingjia Deng and Janyce Wiebe. 2014. Sentiment propagation via implicature constraints. In *EACL*. pages 377–385.
- Haibo Ding and Ellen Riloff. 2016. Acquiring knowledge of affective events from blogs using label propagation. In *AAAI*.
- David K. Elson. 2012. *Modeling Narrative Discourse*. Ph.D. thesis, Columbia University, New York City.
- Song Feng, Jun Seok Kang, Polina Kuznetsova, and Yejin Choi. 2013. Connotation lexicon: A dash of sentiment beneath the surface meaning. In *Association for Computational Linguistics (ACL)*.
- Marie JC Forgeard, Eranda Jayawickreme, Margaret L Kern, and Martin EP Seligman. 2011. Doing the right thing: Measuring wellbeing for public policy. *International Journal of Wellbeing* 1(1).
- Amit Goyal and Ellen Riloff. 2013. A computational model for plot units. *Computational Intelligence* 29(3):466–488.
- Ellen Isaacs, Artie Konrad, Alan Walendowski, Thomas Lennig, Victoria Hollis, and Steve Whittaker. 2013. Echoes from the past: how technology mediated reflection improves well-being. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, pages 1071–1080.
- Paul E Jose, Bee T Lim, and Fred B Bryant. 2012. Does savoring increase happiness? a daily diary study. *The Journal of Positive Psychology* 7(3):176–187.
- Daniel Kahneman. 1999. Objective happiness. *Well-being: The foundations of hedonic psychology* 3:25.
- Daniel Kahneman, Alan B Krueger, David A Schkade, Norbert Schwarz, and Arthur A Stone. 2004. A survey method for characterizing daily life experience: The day reconstruction method. *Science* 306(5702):1776–1780.
- Jun Seok Kang, Song Feng, Leman Akoglu, and Yejin Choi. 2014. Connotationwordnet: Learning connotation over the word+sense network. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Baltimore, Maryland, pages 1544–1554. <http://www.aclweb.org/anthology/P14-1145>.
- Fazel Keshtkar and Diana Inkpen. 2009. Using sentiment orientation features for mood classification in blogs. In *Natural Language Processing and Knowledge Engineering, 2009. NLP-KE 2009. International Conference on*. IEEE, pages 1–6.
- Artie Konrad, Victoria Bellotti, Nicole Crenshaw, Simon Tucker, Les Nelson, Honglu Du, Peter Pirolli, and Steve Whittaker. 2015. Finding the adaptive sweet spot: Balancing compliance and achievement in automated stress reduction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, pages 3829–3838.
- Artie Konrad, Ellen Isaacs, and Steve Whittaker. 2016a. Technology-mediated memory: Is technology altering our memories and interfering with well-being? *ACM Transactions on Computer-Human Interaction (TOCHI)* 23(4):23.
- Artie Konrad, Simon Tucker, John Crane, and Steve Whittaker. 2016b. Technology and reflection: Mood and memory mechanisms for well-being. *Psychology of well-being* 6(1):1–24.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Chris Bizer. 2014. DBpedia - a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web Journal*.
- Peter M Lewinsohn and Christopher S Amenson. 1978. Some relations between pleasant and unpleasant mood-related events and depression. *Journal of abnormal psychology* 87(6):644.
- Peter M Lewinsohn, Robin M Mermelstein, Carolyn Alexander, and Douglas J MacPhillamy. 1985. The unpleasant events schedule: A scale for the measurement of aversive events. *Journal of Clinical Psychology* 41(4):483–498.

- Sonja Lyubomirsky, Laura King, and Ed Diener. 2005. The benefits of frequent positive affect: Does happiness lead to success? *Psychological bulletin* 131(6):803.
- Douglas J MacPhillamy and Peter M Lewinsohn. 1982. The pleasant events schedule: Studies on reliability, validity, and scale intercorrelation. *Journal of Consulting and Clinical Psychology* 50(3):363–380.
- Abraham Harold Maslow. 1943. A theory of human motivation. *Psychological review* 50(4):370.
- Rada Mihalcea and Hugo Liu. 2006. A corpus-based approach to finding happiness. In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*. pages 139–144.
- Gilad Mishne. 2005. Experiments with mood classification in blog posts. In *Proceedings of ACM SIGIR 2005 workshop on stylistic analysis of text for information access*. volume 19, pages 321–327.
- Finn Årup Nielsen. 2011. A new anew: Evaluation of a word list for sentiment analysis in microblogs. *arXiv preprint arXiv:1103.2903*.
- Shereen Oraby, Lena Reed, Ryan Compton, Ellen Riloff, Marilyn Walker, and Steve Whittaker. 2015. And thats a fact: Distinguishing factual and emotional argumentation in online dialogue. *NAACL HLT 2015* page 116.
- Andrew Ortony, Gerald L Clore, and Allan Collins. 1990. *The cognitive structure of emotions*. Cambridge university press.
- J. W. Pennebaker, M. E. Francis, and R. J. Booth. 2001. *LIWC: Linguistic Inquiry and Word Count*.
- Lena Reed, Jiaqi Wu, Shereen Oraby, Pranav Anand, and Marilyn Walker. 2017. Learning lexico-functional patterns for first-person affect. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Kevin Reschke and Pranav Anand. 2011. Extracting contextual evaluativity. In *Proceedings of the Ninth International Conference on Computational Semantics*. Association for Computational Linguistics, pages 370–374.
- Ellen Riloff. 1996. Automatically generating extraction patterns from untagged text. In *Proceedings of the national conference on artificial intelligence*. pages 1044–1049.
- Josef Ruppenhofer and Jasper Brandes. 2015. Extending effect annotation with lexical decomposition. In *6TH WORKSHOP ON COMPUTATIONAL APPROACHES TO SUBJECTIVITY, SENTIMENT AND SOCIAL MEDIA ANALYSIS WASSA 2015*. page 67.
- Irene Russo, Tommaso Caselli, and Carlo Strapparava. 2015. Semeval-2015 task 9: Cliveval implicit polarity of events. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*. Association for Computational Linguistics, Denver, Colorado, pages 443–450. <http://www.aclweb.org/anthology/S15-2077>.
- R.M. Ryan and E.L. Deci. 2000. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American psychologist* 55(1):68.
- Klaus R Scherer, Angela Schorr, and Tom Johnstone. 2001. *Appraisal processes in emotion: Theory, methods, research*. Oxford University Press.
- Klaus R Scherer, Harald G Wallbott, and Angela B Summerfield. 1986. *Experiencing emotion: A cross-cultural study*. Cambridge University Press.
- H Andrew Schwartz, Maarten Sap, Margaret L Kern, Johannes C Eichstaedt, Adam Kapelner, Megha Agrawal, Eduardo Blanco, Lukasz Dziurzynski, Gregory Park, David Stillwell, et al. 2016. Predicting individual well-being through the language of social media. In *Pac Symp Biocomput*. volume 21, pages 516–527.
- Martin EP Seligman, Tayyab Rashid, and Acacia C Parks. 2006. Positive psychotherapy. *American Psychologist* 61(8):774.