

Psychotherapy

Computational Psychotherapy Research: Scaling up the Evaluation of Patient-Provider Interactions

Zac E. Imel, Mark Steyvers, and David C. Atkins

Online First Publication, May 26, 2014. <http://dx.doi.org/10.1037/a0036841>

CITATION

Imel, Z. E., Steyvers, M., & Atkins, D. C. (2014, May 26). Computational Psychotherapy Research: Scaling up the Evaluation of Patient-Provider Interactions. *Psychotherapy*. Advance online publication. <http://dx.doi.org/10.1037/a0036841>

Computational Psychotherapy Research: Scaling up the Evaluation of Patient–Provider Interactions

Zac E. Imel
University of Utah

Mark Steyvers
University of California, Irvine

David C. Atkins
University of Washington

In psychotherapy, the patient–provider interaction contains the treatment’s active ingredients. However, the technology for analyzing the content of this interaction has not fundamentally changed in decades, limiting both the scale and specificity of psychotherapy research. New methods are required to “scale up” to larger evaluation tasks and “drill down” into the raw linguistic data of patient–therapist interactions. In the current article, we demonstrate the utility of statistical text analysis models called topic models for discovering the underlying linguistic structure in psychotherapy. Topic models identify semantic themes (or topics) in a collection of documents (here, transcripts). We used topic models to summarize and visualize 1,553 psychotherapy and drug therapy (i.e., medication management) transcripts. Results showed that topic models identified clinically relevant content, including affective, relational, and intervention related topics. In addition, topic models learned to identify specific types of therapist statements associated with treatment-related codes (e.g., different treatment approaches, patient–therapist discussions about the therapeutic relationship). Visualizations of semantic similarity across sessions indicate that topic models identify content that discriminates between broad classes of therapy (e.g., cognitive–behavioral therapy vs. psychodynamic therapy). Finally, predictive modeling demonstrated that topic model-derived features can classify therapy type with a high degree of accuracy. Computational psychotherapy research has the potential to scale up the study of psychotherapy to thousands of sessions at a time. We conclude by discussing the implications of computational methods such as topic models for the future of psychotherapy research and practice.

Keywords: psychotherapy, topic models, linguistics

I believe that some aspects of psychoanalytic theory are not presently researchable because the intermediate technology required . . . does not exist. I mean auxiliaries and methods such as a souped-up, highly developed science of psycholinguistics, and the kind of mathematics that is needed to conduct a rigorous but clinically sensitive and psychoanalytically realistic job of theme tracing in the analytic protocol (Meehl, 1978, p. 830).

Advances in technology have revolutionized research in much of psychology and health care, including major developments in pharmacology, neuroscience, and genetics. Yet, the science of patient–therapist interactions—the core of psychotherapy process research—has remained fundamentally unchanged for 70 years. Patients fill out surveys, or human coders rate some aspect of the interaction. Thus, although psychiatric and psychological guidelines recommend psychotherapy as a first line treatment for a number of mental disorders (American Psychiatric Association, 2006), we still know relatively little about how psychotherapy works. As Meehl noted, existing research methods remain limited in their ability to explore the structure of verbal exchanges that are the essence of most psychotherapy. In the current article, we move toward an answer to Meehl’s request for a “souped up mathematics” to mine the raw linguistic data of psychotherapy interactions. In traditional research on psychotherapy, human judgment and related behavioral coding are the rate-limiting factor. In this article, we introduce a computational approach to psychotherapy research that is informed by traditional methods (e.g., behavioral coding) but does not rely on them as the primary data source. The key innovation in this computational approach is drawing on methods from computer science and machine learning that allow the direct, statistical analysis of session content, scaling up research to thousands of sessions.

Zac E. Imel, Department of Educational Psychology, University of Utah; Mark Steyvers, Department of Cognitive Sciences, University of California, Irvine; David C. Atkins, Department of Psychiatry and Behavioral Sciences, University of Washington.

Funding for the preparation of this article was provided by National Institute of Drug Abuse (NIDA) of the National Institutes of Health under award R34/DA034860, the National Institute on Alcohol Abuse and Alcoholism (NIAAA) under award R01/AA018673, and a special initiative grant from the College of Education, University of Utah. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or University of Utah. The authors also thank Alexander Street Press for consultation related to the analysis of psychotherapy transcripts.

Correspondence concerning this article should be addressed to Zac E. Imel, Department of Educational Psychology, University of Utah, 1705 Campus Center Drive, Room 327, Salt Lake City, UT 84112. E-mail: zac.imel@utah.edu

Many Distinctions, but Is There a Difference?

Some estimates indicate that there are >400 different name brand psychotherapies (Lambert, 2013); each treatment offers a different approach to helping patients with psychological distress. Although the clinical rationales and approaches differ, it is not clear that actual practices of these psychotherapies are meaningfully distinct. Potential differences in the process and outcome of psychotherapies have been a focus of psychotherapy science for over a century. As a comparison, there are many different drug therapies. However, the unique ingredients of treatments are chemical (and patentable). Thus, the actual distinctiveness of treatments is known, even if the specific mechanism of action or relative efficacy is not. In psychotherapy, the treatment consists primarily of words, and although cognitive-behavioral oriented treatments (CBT) might focus strongly on patient behavior, the treatment is still verbally mediated (Wampold, 2007). Accordingly, scientific classification of treatments is more nebulous. What is considered a “taxon” of cognitive-behavioral therapy may vary widely across experts and practitioners, with some definitions so broad as to include any scientifically justifiable intervention and others restricted to very specific psychological mechanisms (see Baardseth et al., 2013). This ambiguity is quite old, reaching back to debates between Freud and his early followers and can be found in current research comparing various cognitive-behavioral psychotherapies and modern variants of psychoanalysis (e.g., psychodynamic psychotherapy; Leichsenring et al., 2013).

Some have argued that differences between psychotherapies are cosmetic (like the difference between generic ibuprofen and Advil) and that the underlying mechanisms of action are common across different approaches (Wampold, 2001). Meta-analyses generally suggest that most treatment approaches are of comparable efficacy (e.g., Benish, Imel, & Wampold, 2008; Imel, Wampold, Miller, & Fleming, 2008), and process studies cast doubt on the relationship between treatment-specific therapist behaviors and patient outcomes (Webb, DeRubeis, & Barber, 2010). One leading addiction researcher commented that, “. . . there is little evidence that treatments work as purported, suggesting that as of yet, we don’t know much about how brand name therapies work” (Morgenstern & McKay, 2007, p. 87S). Are the 400 psychotherapies we have today unique, medical treatments? Or, are the different psychotherapies largely similar, distinguished by packaging that obscures what are mostly common components?

Given that psychotherapy is a conversation between patient and provider, the distinctiveness of a therapy approach should be found in the words patients and therapists use during their sessions. Yet, this is precisely where we find a fundamental methodological gap in psychotherapy research. The source data and information are linguistic and semantic, but the available tools used to study psychotherapy are not. Research on the active ingredients of psychotherapy has primarily relied on patient or therapist self-report measures (e.g., see reviews of empathy and alliance literature; Elliott, Bohart, Watson, & Greenberg, 2011; Horvath, Del Re, Flückiger, & Symonds, 2011) or on behavioral coding systems, wherein human “coders” make ratings from audio or video recordings of the intervention session according to a priori theory-specific criteria (Crits-Christoph, Gibbons, & Mukherjee, 2013).

Attempts at behavioral coding have varied in their depth from general, topographical assessments of the session such as those

used in many cognitive-behavioral treatments (e.g., did the therapist ask about homework or set an agenda?) to highly detailed utterance-level coding systems (e.g., verbal response modes, Stiles, Shapiro, & Firth-Cozens, 1988; Motivational Interviewing Skills Code, Moyers, Miller, & Hendrickson, 2005). However, behavioral coding as a technology has not fundamentally changed since Carl Roger’s first recorded a psychotherapy session in the 1940s (Kirschenbaum, 2004), and coding carries a number of disadvantages. It is extremely time consuming, and reliability can be problematic to establish and maintain. In addition, there is no potential for human coding to scale up to larger applications (i.e., coding 1,000 sessions takes 1,000 times longer than coding 1 session; thus, monitoring the quality of psychotherapy in a large scale naturalistic setting is not feasible over time). There is little flexibility—coding systems only code what they code. They must be developed a priori and cannot discover new meaning not specified in advance by the researcher. More substantively, coding systems are by nature extremely reductionistic—reducing the highly complex structure of natural human dialogue to a small number of behavioral codes.

Given these limitations, it is not surprising that the vast majority of raw data from psychotherapy is never analyzed and questions central to psychotherapy science remain either unanswered or impractical to address. Most content analyses of what patients and therapists actually discuss in psychotherapy are restricted to qualitative efforts that can be rich in content but by their nature are small in scope (e.g., Greenberg & Newman, 1996). Although qualitative work remains important, the labor intensiveness of closely reading session content means that the vast majority of psychotherapy data are never analyzed. Consequently, the majority of psychotherapy studies are published without any detail as to what the specific conversations between patients and therapists actually entailed. Beyond the general theoretical description of the treatment outlined in manuals, what did the patients and therapists actually say? Are the different psychotherapies we have today linguistically unique? Or, do therapists who provide different name brand therapies say largely similar things? What specific therapist interventions, and in what combination are most predictive of good versus bad outcomes? These basic questions form the backdrop of every therapist’s work, but have been impractical to consider given the current technology of behavioral coding and qualitative analysis.

A critical task for the next generation of psychotherapy research is to move beyond the use of behavioral coding to mine the raw verbal exchanges that are the core of psychotherapy, including acoustic and semantic content of what is said by patients and therapists. The use of discovery-oriented machine learning procedures offer new ways of exploring and categorizing psychotherapies based on the actual text of the patient and therapist speech.

Text Mining and Psychotherapy

The amount of data generated every day (e.g., digitized books, e-mail, video, newspapers, blog posts, twitter, electronic medical records, cell phone calls) has expanded exponentially in the last decade, with implications for business, government, science, and the humanities (Hilbert & Lopez, 2011). Developments in data-mining procedures have revolutionized our ability to analyze and understand this vast amount information, particularly in the area of

text—sometimes called “computational linguistics” or “statistical text classification” (Manning & Schütze, 1999). Google books “n-gram” server (<https://books.google.com/ngrams>) allows for the evaluation of trends in single words (i.e., unigrams) or word combinations (bigrams, trigrams) in books. A recent article analyzed words in 4% of all books (5,195,769 volumes), showing that patterns of emotion word use tracked in expected directions with major historical events (e.g., a sad peak during World War II; Acerbi, Lampos, Garnett, & Bentley, 2013).

There is a small literature that demonstrates the utility of computational linguistic approaches for the analysis of psychotherapy data. The majority of these studies rely on human-defined computerized dictionaries in which a software program classifies words or sets of words into predefined categories. In an early study Reynes, Martindale, and Dahl (1984) found that “linguistic diversity” was higher in more productive sessions. In addition, Mergenthaler and his colleagues have published several studies demonstrating that emotion and abstraction word usage discriminates between improved and unimproved cases (e.g., Mergenthaler, 2008; see also Anderson, Bein, Pinnell, & Strupp, 1999). Studies that have used dictionary-based strategies hold promise, but also have important limitations. First, perhaps because large corpora of psychotherapy transcripts are hard to find, these studies have generally been limited in scope ($n < 100$), reducing the value added of a computerized technology that can evaluate a large set of sessions (i.e., 1,000 or 10,000) in a short amount of time. Second, computerized dictionaries are limited by the categories created by humans—the computer cannot “learn” new categories.

Finally, dictionaries cannot generally accommodate the effect of context on semantic meaning (e.g., “dark” may reference a mood or the sky at night).

Topic Models

One specific text-mining approach that holds promise for psychotherapy transcript data are *topic models* (also called, Latent Dirichlet Allocation; Blei, Ng, & Jordan, 2003). Topic models are data-driven, machine learning procedures that seek to identify semantic similarity among groups of words. Similar to factor analysis in which observed item values are functions of underlying dimensions, topic models view the observed words in a passage of text as a mixture of underlying semantic topics. An advantage of topic models is that they construct a linguistic structure from a set of documents inductively, requiring no external input, but can also be used in a supervised fashion to learn semantic content associated with particular codes or metadata (where metadata is any data outside of the text itself; Steyvers & Griffiths, 2007).

There is recent work using these models to explore the structure of National Institute of Health grant applications (Talley et al., 2011), publications from the Proceedings of the National Academy of Sciences (Griffiths & Steyvers, 2004), articles from the New York Times (Rubin, Chambers, Smyth, & Steyvers, 2012), and the identity of scientific authors (Rosen-Zvi, Chemudugunta, Griffiths, Smyth, & Steyvers, 2010). Perhaps more strikingly, topic models have been used in the humanities to facilitate “distant reading” in comparative literature such that hypotheses in literary criticism can be tested vis-à-vis the entire corpus of relevant work (e.g., exploring stylistic similarities in poems, see Kao & Jurafsky, 2012; Kaplan & Blei, 2007).

With a few exceptions, topic models have yet to be applied to psychotherapy data (see Atkins et al., 2012 and also Salvatore, Gennaro, Auletta, Tonti, & Nitti, 2012 who used a derivative of latent semantic analysis—a forerunner to topic models; Landauer & Dumais, 1997). However, similar to the news articles, novels, and poems noted above, the words used during psychotherapy sessions by patients and therapists can be viewed as a large collection of text with a complex topical structure. The number of words generated during psychotherapy is quite large. A brief course of psychotherapy for a given patient may consist of 5–10 hr of unstructured dialogue including 12,000–15,000 words per hour (approximately 60,000 to 150,000 words, longer courses of treatment >1 million words). In 2011, a PubMed search revealed 932 citations for psychotherapy clinical trials (out of 10,698 across all years). As a conservative estimate, if we consider 500 studies per year, 50 participants per study, 5 sessions per participant, and 10,000 words per session, this leads to an estimate of 125 M words of psychotherapy text per year from clinical trials alone. Regardless of the specific estimate, it is clear that a huge amount of psychotherapy data is generated every year and that this number is likely to increase. The use of discovery-oriented text-mining procedures such as topic models could facilitate new ways of exploring and categorizing psychotherapies based on the actual content of the patient and therapist speech (rather than labels established by schools of psychotherapy).

Current Study

To evaluate the potential of topic models to “learn” the language of psychotherapy, we applied two different types of topic models to transcripts from 1,553 psychotherapy and psychiatric medication management sessions. Our first goal was to verify that topic models would estimate clinically relevant semantic content in our corpus of therapy transcripts. Second, we determined if semisupervised models could identify semantically distinctive content from different treatment approaches and interventions (e.g., therapist “here and now” process comments about the therapeutic relationship within a session). A third aim was to explore the overall linguistic similarity and distinctiveness of sessions from different treatment types (e.g., psychodynamic vs. humanistic/experiential). Our final goal was to classify treatment types of new psychotherapy sessions automatically, using only the words used during the session.

Method

Data Sources

The data for the current proposal come from two different sources: (a) a large, general psychotherapy corpus that includes sessions from a diverse array of therapies, and (b) a set of transcripts focused on motivational interviewing (MI), a specific form of cognitive-behavioral psychotherapy for alcohol and substance abuse.

General psychotherapy corpus. The general corpus holds 1,398 psychotherapy and drug therapy (i.e., medication management) transcripts (approximately 2.0 million talk turns, 8.3 million word tokens including punctuation) pulled from multiple theoret-

ical approaches (e.g., cognitive behavioral; psychoanalysis; MI; brief relational therapy).

The corpus is maintained and updated by the “Alexander Street Press” (<http://alexanderstreet.com/>) and made available via library subscription. In addition to transcripts, there is associated metadata such as patient ID, therapist ID, limited demographics, session numbers when there was more than a single session, therapeutic approach, patient’s primary symptoms, and a list of subjects discussed in the session.

The list of symptoms and subjects was assigned by publication staff to each transcript, and no interrater reliability statistics were available. All labels were derived from the *DSM-IV* and other primary psychology/psychiatry texts. Many sessions were conducted by prominent psychotherapists who developed particular treatment approaches (e.g., James Bugental, existential; Albert Ellis, rational emotive; Carl Rogers, person-centered; William Miller, MI), and hence may serve as exemplars of these treatment approaches. To facilitate analysis we categorized each psychotherapy session into one of five treatment categories, (a) Psychodynamic (e.g., psychoanalysis, brief relational therapy, psychoanalytic psychotherapy), (b) Cognitive behavioral therapy (e.g., rationale emotive behavior therapy, MI, relaxation training, etc.), (c) Experiential/Humanistic (e.g., person-centered, existential), (d) other (e.g., Adlerian, Reality Therapy, Solution-Focused, as well as group, family, and marital therapies), and finally (e) Drug therapy or medication management. However, in some cases, when a label was missing or more than one treatment label was assigned to a session, collateral information in the metadata was used to assign a single specific treatment label (i.e., a well known therapist associated with a specific intervention, reported use of specific interventions, and/or inspection of the raw transcript). If there was no collateral information or an appropriate label could not be determined, the first listed intervention was chosen as the treatment name or the treatment label and category was left missing. In addition to treatment category, analyses used one subject label, “counselor–client relations.” This session-level label (i.e., applied to an entire session) was assigned to a transcript when there was a discussion about the patient–therapist relationship or interaction during the therapy.

MI corpus. We supplemented the general corpus above with a set of MI sessions (Miller & Rollnick, 2002; $n = 148$, 30,000 talk turns, 1.0 million word tokens). Transcripts are a subset of sessions from five randomized trials of MI for drug or alcohol problems, including problematic drinking in college freshman (Tollison, Lee, Neighbors, & Neil, 2008), 21st birthdays and spring break (Neighbors et al., 2012), problematic marijuana use (Lee et al., 2014), and drug use in a public safety-net hospital (Krupski et al., 2012). Each study involved one or more in-person treatment arms that received a single session of MI. Sessions were transcribed as part of ongoing research focused on applying text-mining and speech signal-processing methods to MI sessions (see, e.g., Atkins, Steyvers, Imel, & Smyth, in press).

Data Analysis

The linguistic representation in our analysis consisted of the set of words in each talk turn. A part-of-speech tagger (Tou-

anova, Klein, Manning, & Singer, 2003) was used to analyze the types of words in each talk turn. We kept all nouns, adjectives, and verbs and filtered out a number of word classes such as determiners and conjunctions (e.g., “the,” “a”) as well as pronouns. The resulting corpus dramatically reduces the size of the corpus to 1.2M individual words across 223K talk turns. We applied a topic model with 200 topics to this data set, treating each talk turn (either patient or therapist) as a “document.” In the topic modeling literature, the document defines the level at which words with similar themes are grouped together in the raw data. We could define documents in a number of ways (e.g., all words in the session or all words from a specific person), but we have found in previous research within clinical psychology (Atkins et al., 2012) that defining documents by talk turns enhances the interpretability of the resulting topics. In a topic model, each topic is modeled as a probability distribution over words and each document (talk turn) is treated as a mixture over topics. Each topic tends to cluster together words with similar meaning and usage patterns across talk turns. The probability distribution over topics in each talk turn gives an indication of which semantic themes are most prevalent in the talk turn. For further details on topic models, see Atkins et al. (2012).

Results

Exploration of Specific Topics

First, we used topic models to explore what therapists and patients talk about. As noted earlier, topic models estimate underlying dimensions in text, which ideally capture semantically similar content (i.e., the underlying “topics”). Thus, in applying topic models to psychotherapy transcripts, an initial question is whether the models extract relevant semantic content? Table 1 presents 20 selected topics (of 200 total) from an unsupervised topic model applied to all session transcripts (i.e., these topics were generated inductively without any input from the researchers). It is clear that the words in each topic provide semantically related content and capture aspects of the clinical encounter that we might expect therapists and patients to discuss. We have organized topics into four areas—(a) Emotions/Symptoms, (b) Relationships, (c) Treatment, and (d) Miscellaneous. Similar to factor analysis, all labels were supplied by the current authors—the model itself simply numbers them. The top 10 most probable words for each topic are provided along with author-generated topic labels to aid interpretation. For example, the emotion category includes several symptom-relevant topics. Topic 15 (Depression) includes many of the specific symptom criteria for depression (e.g., *sadness*, *energy*, *hopelessness*; the word “depression” is the 16th most probable word), and topic 149 (Anxiety) includes words relevant to the discussion of a panic attack.

The relationship category illustrates how a topic model can handle differences in meaning depending on context. Topic 146 (Sex) and 60 (Intimacy) include derivatives of the words *relationship* and *sex*. In Topic 60, these words occur in the context of words such as *closeness*, *intimacy*, *connection*, and *open*, suggesting these words had different implications than when they occur in Topic 60, which includes words such as *desire*,

Table 1
Selected Topics

Emotions				
#149: Anxiety anxiety, nervous, anxious, panic, attack, attacks, tense, calm, depressed, hyper	#124: Crying crying, cry, hurt, cried, upset, emotional, tears, face, start, sudden	#156: Hurt feelings feelings, hurt, strong, emotions, express, emotion, intense, touch, emotional, hurts	#100: Enjoyment enjoy, fun, excited, enjoying, find, enjoyed, pleasure, exciting, interest, company	#15: Depression self, fine, low, sad, appetite, hopeless, helpless, esteem, irritable, energy
Relationships				
#146: Sex sex, sexual, normal, relationship, healthy, desire, satisfied, involved, marriage, enjoy	#168: Pregnancy baby, boy, pregnant, child, born, boys, girl, son, mother, age	#73: Conflict hate, fight, stand, awful, horrible, fighting, terrible, argument, argue, hated,	#76: Family roles sister, brother, older, younger, mother, family, daughter, father, mom, sisters	#60: Intimacy relationship, relationships, close, sexual, involved, develop, intimate, connection, open, physical
Treatment				
#196: Medication wellbutrin, prozac, zoloft, medicine, medicines, effexor, lexapro, add, generic, lamictal	#198: Behav pattern behavior, pattern, aggressive, least, example, irrational, personality, conscious, follow, identify	#135: MI survey information, questions, feedback, helpful, survey based, interested, great, useful, use	#131: Goal setting set goal, goals, expectations, successful, success, own, setting, working, accomplish	#69: Subst Use Tx treatment, program, need, options, stay, meetings, available, sound, use, option
Miscellaneous				
#65: Change difference, noticed, notice, big, huge, change, improvement, happens, significant, differences	#80: Medical doctor, hospital, cancer, doctors, disease, nurse, surgery, sick, patients, medical	#104: Drinking alcohol, drinking, social, effects, tired, outgoing, drunk, situations, sounds, relaxed	#120: Appearance wear, hair, clothes, looking, feet, dress, stand, wearing, shoes, ugly	#50: Acceptance accept, find, change, willing, least, accepting, accepted, situation, hope, possibility

enjoy, and *satisfied*. The basic topic model can infer differing meaning of identical words (e.g., play used in reference to theater vs. children) as long as the documents that the words occur in have additional semantic information that would inform the distinction (Griffiths, Steyvers, & Tenenbaum, 2007). In the treatment category, topic 196 includes a number of medication names and is clearly related to discussions of psychopharmacological treatment. Topic 198 (Behavior Patterns) includes words that might be typical in the examination of behavior/thought patterns (e.g., *irrational*, *pattern*, *behavior*, *identify*). We considered labeling this topic “CBT,” given words that might be found in an examination of thoughts in cognitive therapy. However, we found that this topic was actually more prevalent in psychodynamic sessions as compared with CBT sessions. This finding highlights the complexities of topic models. While the model returns a cluster of words, the researcher must infer what the cluster means.

Identification of Therapist Interventions

To demonstrate the utility of a topic model in the discovery of language specific to different approaches to psychotherapy, we used a “labeled” topic model (Rubin et al., 2012) wherein the model learns language that is associated with a particular label—in the present case a session-level label that identifies the type of psychotherapy (e.g., CBT vs. Psychodynamic).

We used the output from this model to identify specific therapist talk turns that were statistically representative of a given label. In the general psychotherapy corpus, there were no labels or codes for

talk turns, only for the session as a whole. Given the labels for each session and the heterogeneity of word usage across sessions, the model “learns” which talk turns were most likely to give rise to a particular label for the entire session.

In Table 2, we provide four highly probable talk turns for six different treatments. The depicted statements are what might be considered prototypic therapist utterances for each treatment. Client-centered talk turns appear to be reflective in nature, while utterances in rationale emotive behavior therapy have a quality of identifying irrational thought patterns. Brief relational interventions focus on here and now experiences, and the selected talk turns for MI were those typical for the brief structured feedback session that therapists were trained to provide in several of the MI clinical trials included in the corpus.

Table 2 presents results from a labeled topic model using psychotherapy type as the label categorizing a session. We explored whether the model could learn more nuanced, psychological labels, focusing on “client–counselor relations”—a code that was used to label sessions that included discussions between client and therapist about their relationship/interaction. As with the identification of therapist talk turns, the client–counselor relations code was assigned to an entire transcript. Consequently, the model must learn to discriminate between language in these sessions that is irrelevant to the label (e.g., general questions, scheduling, pleasantries, other interventions, etc.) and language that involves the client and therapist talking about their relationship. Table 3 provides the five most probable therapist talk turns associated with the client–counselor relations label. Each talk turn is clearly related to

Table 2
Most Probable Talk Turns for Specific Treatment Labels

Treatment label	Example therapist talk turns assigned by model
Drug therapy	No trouble getting to sleep or staying asleep? and how's your energy level holding up, you doing okay? So, so in this, lorazepam. So in this sweep or time, over these 3 months, how's your mood been? Separate from this setback, how's your mood been in general? I saw you, ah, . . . your mood was okay. You were recently stable. You were not sad or anxious or irritable. Your appetite and sleep were fine. Your energy was good. You were exercising, some revving up in fall. Again, I had not seen you since August before. So and August was all the drama over the breakup. I'll give you the 25s. So let me write for the 200s. Okay, lamictal 200 and I'll write for Wellbutrin, 3 of the 150s?
Client-centered therapy	And its kind of like, I guess it's like it felt great to be to finally sleep with [Name]. Its kind of like, like you really see yourself kind of extending yourself. Yeah, I kind of sense like you really feel like a blank today. Yeah. It's like you really kind of feel like emotionally you may . . . you're deadened inside or keep yourself guarded.
Psychoanalysis	I see. It's as though you were being a kind of medium for us. So you're, you're afraid on the one hand to let your thoughts go to something else because you feel you're leaving the subject, eh. I gather from what you say that you must wonder whether or not to tell your parents that you've started analysis. Well, I don't know that there is. It just strikes me that you tell me the dream, you don't say anything about it, and then say what you've been thinking about now, among other things, whether you should learn to use a diaphragm, and whether you need it and why do you need it. Which are the questions in the dream.
Brief relational therapy	But how did it feel? Did it feel like I was letting you down? or did it feel like I was wimping out? did it? And I was sort of trying to explore what was going on between you and me. I'm asking you now how does it feel to say it? Let us talk about what is going on here? How are you feeling right now?
Rational emotive therapy	Ok. And what about the bigger one, that "[Name] should do things the way that I want them done and if he doesn't, he's an asshole"? You are like a star student they were like whose. Where was I? So it is more of the catastrophic thinking and there is some self-doubting as well? Well maybe there is three, because is there the 'I can't stand that he hasn't called. I can't stand this.'? . . . To succeed. That is kind of your main or irrational belief. "I should not have to work as hard as other people to succeed."
Motivational interviewing	Okay final two continue to minimize my negative impacts on the environment. How if at all does Marijuana use affect attainment of that goal? Mm-hmm and how might that fit into your plans for spring break. Okay so eight drinks over two hours would put you at a point one seven two. So this next part is about BAC or Blood Alcohol Content.

Note. The four most representative therapist talk turns for five specific treatments. Direct quotations from session transcripts reproduced with permission by Alexander Street Press (<http://alexanderstreet.com/>).

a therapist making a comment about the patient–therapist interaction.

Discrimination of Treatment Approaches

In addition to low-level identification of therapist statements, we used topic models to make high-level comparisons related to the linguistic similarity of sessions. How similar are sessions, given the semantic content identified by the topic model? We used the output from the unsupervised topic model to explore the semantic similarity of 1,318 sessions across four treatment categories (i.e., Medication Management, Psychodynamic, CBT, Humanistic/Existential). Specifically, it is possible to assign individual words within sessions to one of the 200 topics. The sum of the words in each topic for each session provide a session-level summary of the session's semantic content—a model-based score on each of 200 topics for each of the 1,318 sessions.¹ Given these semantic summaries of each session, we then computed a correlation matrix of each session with every other session. A high correlation between two sessions indicates similar semantic content, defined by the 200 topics of the topic model. Because a $1,318 \times 1,318$ matrix of correlations would be utterly unreadable, we present the corre-

lation matrix visually using color-encoded values for the correlations.

This style of visualization is referred to as a heatmap, as the initial versions often used red to yellow coloring to note the intensity of the numeric values. In Figure 1, the color scale on the right shows how correlation values are mapped to specific colors: Orange and red pixels represent highly correlated sessions, and blue and green pixels indicate little correlation in topic frequencies. The correlation matrix was purposefully organized by treatment category. We have highlighted several highly correlated blocks of sessions that represent (a) a set of highly structured MI feedback sessions from a clinical trial, (b) a large number of sessions from a single case of psychoanalysis, and (c) several sessions from a single case of client-centered therapy. Sessions within treatment category are generally more correlated than outside of category (e.g., medication management sessions generally have similar topic loadings that are heavily driven by drug names, dosing schedules, etc.). However, correlations across psychody-

¹ The scores were also divided through by total number of words per session so that sessions with different lengths did not skew the results.

Table 3

Example Therapist Talk Turns Assigned by Model for Label “Client–Counselor Relations”

I am asking you questions. I am asking you questions and asking you to look at stuff and you are joking and giggling again.
 I guess I could try to explain it again. I'm just wondering if any explanation I give because we—we have—we have discussed what we're doing in therapy or how this works.
 Well you might garner sort of what it feels like just to be able to when I do different things. How it makes you feel that we bring attention to it sometimes. And—and your reactions to it are really important, 'cause in the outside world, your reactions are going to be telling you what your experience is.
 Okay, so let me come back for a second. Because what you are talking about is important and it is a big part of what this impasse that we have been having is all about. I was curious and I am not sure if you answered about the laughing today.
 Well no . . . wait. There is something. We were on the cusp of discussing something really important when this came up? Let me ask you the question more directly. Did you want to discuss this whole thing with [Name] in the session?

Note. The five most probable talk turns for the client-counselor relations label. Direct quotations from session transcripts reproduced with permission by Alexander Street Press, Counseling and Psychotherapy Transcripts, Client Narratives, and Reference Works (<http://alexanderstreet.com/>).

namic and humanistic/experiential session were often moderate such that it is difficult to separate them from visual inspection of the plot. In addition, there are pockets of sessions that are correlated across categories. For example, the zoomed in portion of the heatmap depicted in the lower right portion of Figure 1 highlights several psychodynamic and cognitive–behavioral sessions that had very similar topic loadings. Interestingly, several of these sessions had both CBT and brief relational therapy labels, suggesting that the model was sensitive to potential overlap in content that was identified by the human raters who created the database.

Figure 2 is an alternative visual representation that highlights the semantic similarities and differences across sessions, called a multidimensional scaling (MDS; Cox & Cox, 2000) plot. Using the same session-level topic scores from the correlation matrix above, MDS treats each session's 200 values as a set of coordinates (in a 200 dimension, mathematical space).

Thus, the topic model-based semantic scoring can be used to define distance values of each session from every other session within a 200 dimension semantic space. Somewhat similar to factor analysis, MDS finds an optimal, lower dimensional space that best represents the overall distance matrix; Figure 2 plots the results of the MDS. Each color-coded dot represents a single session. There was separation between treatment types such that treatment classes were broadly grouped together. However, there was variability within treatment approaches. For example, one set of CBT sessions (denoted in red) are notably different from other sessions. These are the structured MI sessions that all focus on drug or alcohol problems. Other CBT sessions are much more similar to other treatment approaches, and interestingly, appear to lie in between the highly structured medication management sessions and much less structured experiential sessions. In addition, we highlighted one medication management session that was distinct from the other medication management sessions, located much closer to experiential psychotherapy sessions.

An inspection of this transcript revealed that there was no direct discussion of medications or dosage, potentially indicating a medication provider who focused on providing psychotherapy rather than checking medication dosage and side effects.

The previous results are exploratory visualizations demonstrating how semantic content from a topic model could distinguish categories of psychotherapy. Our final analysis examined how accurately the 200 topics could discriminate these four classes of psychotherapy sessions, using a type of multinomial logistic regression. We used a machine learning regression model called a

random forest model using the 200 topics as predictors (Breiman, 2001). Random forest models are a type of ensemble learner, in which many regressions are fit simultaneously and then aggregated into a single, overall prediction model.² The prediction accuracy of the model is tested using sessions that were *not* used during the training phase. This is a type of cross-validation in which the prediction accuracy of a model is tested on data points that were not included in the model creation. The overall, cross-validated classification error rate was 13.3%, showing strong predictive ability of the topic model-based predictors. As we saw in the earlier visualizations, the semantic information identified by the topic model is highly discriminative of the classes of psychotherapy. Table 4 shows the specific types of errors that the model makes (called a confusion matrix). The rows contain the true psychotherapy categories, and the columns have the model predictions. The counts along the main diagonal indicate correct classifications by the model and off-diagonal elements are errors. Not surprisingly, the model is most accurate at identifying medication management sessions but is also quite accurate with experiential psychotherapy. It is less accurate with CBT and Psychodynamic sessions, which are more likely to be confused as experiential psychotherapy. This makes clinical sense, as the hallmarks of good experiential psychotherapy are reflective listening skills, which are common (though not as strongly emphasized) to CBT and Psychodynamic treatments.

Discussion

We used a specific computational method, topic models, to explore the linguistic structure of psychotherapy. Without any user input, these models discovered sensible topics representing the issues that therapists and patients discuss, and facilitated a high-level representation of the linguistic similarity of sessions wherein we could identify specific cases, potentially overlapping content across treatment approaches, as well as outlier sessions. By including human-generated session labels, topic models learned therapist statements associated with different treatment approaches and interventions, including ther-

² For the present analyses, we created 2,000 new datasets, each with 1,318 sessions sampled with replacement from the original sessions. Next, on each of the 2,000 samples a classification and regression tree model is fit, but only using a subset of the total predictors. Thirty predictors were selected randomly within each bootstrap-generated dataset. This process results in 2,000 sets of regression results, which are then combined into an overall prediction equation.

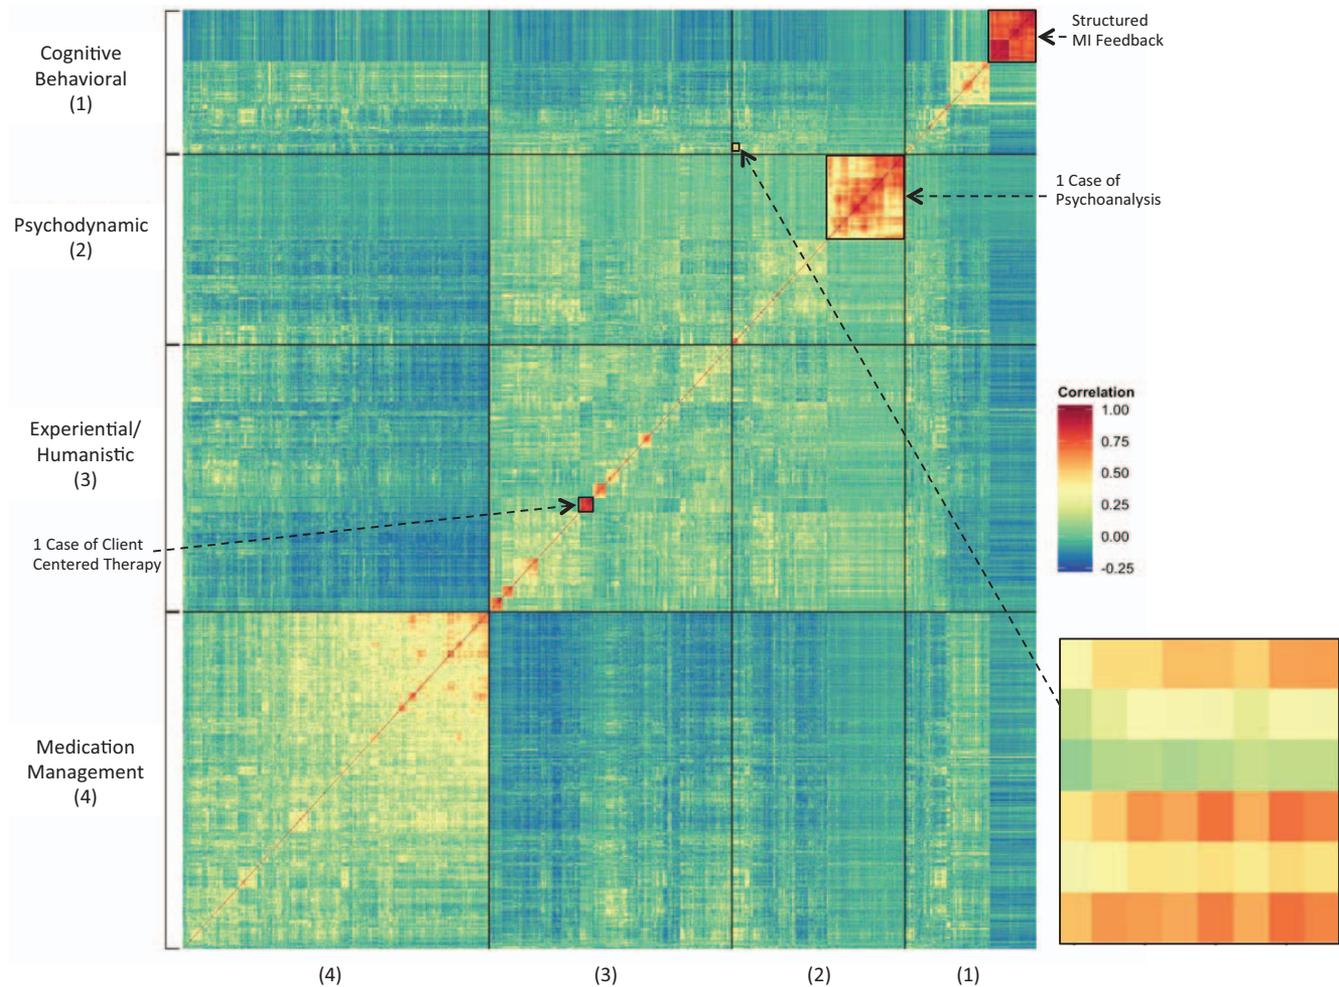


Figure 1. A $1,318 \times 1,318$ heatmap, depicting the correlation of topics across each session. The color scale on the right shows how correlation values are mapped to specific colors. The correlation matrix is organized by treatment category and several select groupings of sessions are highlighted. The color version of this figure appears in the online article only.

apist comments about therapeutic relationship, which are often considered among the more complex interventions in the therapist repertoire. Using only the words spoken by patients and therapists, the topic model classified treatment sessions with a high degree of accuracy.

Limitations

Although the present study represents—what we believe is—the largest comparative study of linguistic content from psychotherapy ever conducted, there are important limitations that we will discuss prior to highlighting potential implications. First, in terms of the data, the combined general psychotherapy and MI corpus is very heterogeneous along several dimensions (e.g., treatment approach, topics of discussion, etc.), but it is certainly not a random sample of general psychotherapy and they were not necessarily collected for research purposes. Although the diversity of the corpus facilitates the examination of differences between approaches, the database is also highly unbalanced. There is an overrepresentation of select cases (>200

sessions from one case), and relatively few sessions from many approaches. For example, CBT is relatively under-represented relative to its empirical standing in modern psychotherapy research, and much of the CBT are MI sessions that may not be representative of other more modal CBT interventions (e.g., prolonged exposure, cognitive therapy for depression). As a result linguistic differences between treatments may be confounded with other differences in the selected sessions not related to approach (i.e., therapists, symptoms, idiosyncratic patient factors, etc.). The labeling of sessions was not done with standard adherence manuals, such that no estimates of reliability are possible. There is no symptom severity or diagnostic data beyond session-level labels that indicate that depression was discussed in a session. There is no audio, which is clearly important to the evaluation of psychotherapy.

The model itself contains a number of important limitations. First, the topic model we used did not include information regarding the temporal ordering of words and talk turns. This is common to most topic models, which make a “bag of words” assumption that word

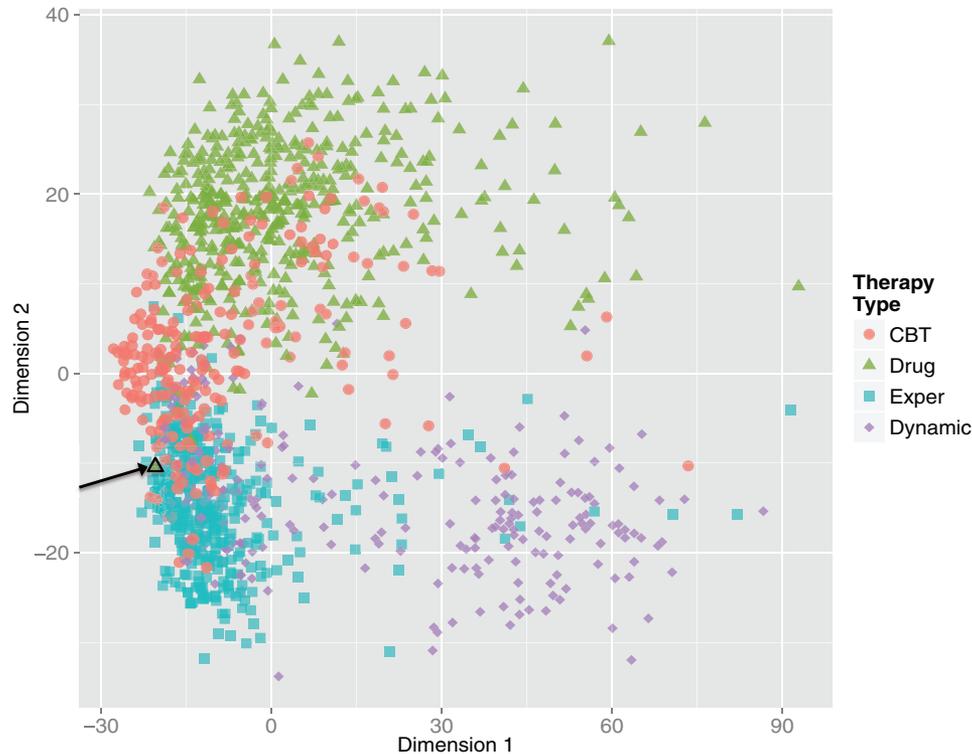


Figure 2. Multidimensional scaling of 1,318 sessions in a 200 topic space. Colors correspond to different treatment approaches. One outlier medication management session is circled in black. The color version of this figure appears in the online article only.

order is not critical. For most prior applications (e.g., news articles and scientific abstracts), this may be a reasonable assumption, but for spoken language it is clearly quite tenuous. In addition, although the removal of specific words like pronouns reduces the complexity of the data, it is likely that these words are quite important in psychotherapy and general human interactions (Williams-Baucom, Atkins, Sevier, Eldridge, & Christensen, 2010). The model was also restricted to text and did not have access to the acoustic aspects of these treatment interactions, which are also important (Imel et al., 2014). Future studies should incorporate the above features.

Table 4
Confusion Matrix of True Versus Predicted Treatment Labels

True category	Predicted category				Class error
	CBT	Drug	Exper	Dynamic	
CBT	153	17	42	8	.30
Drug	4	454	7	1	.03
Exper	0	6	351	12	.05
Dynamic	4	8	66	185	.30

Note. CBT = cognitive behavioral therapy; Drug = medication management; Exper = experiential/humanistic therapy; Dynamic = psychodynamic or psychoanalytic therapy. The bolded, diagonal elements represent correct classifications by the model, and off-diagonal elements represent errors. The final column has the classification errors of the model for each category of therapy (i.e., row). The overall error-rate is .13.

Transcription is a limitation of expanding this work. To use these methods, researchers would be required to transcribe thousands of sessions from clinical trials. While this is an important practical limitation, we believe the primary reason that transcription remains uncommon is that the methods available to analyze transcript data in psychotherapy are labor intensive. In comparison with the cost of a clinical trial, the cost of basic transcription is minimal and could proceed in parallel to the clinical trial. Thus, although transcription would add costs to clinical trials, the costs would be trivial compared with the potential long-term scientific impact of retaining the raw ingredients that were involved in the change process. It is also important to note that automated speech recognition techniques continue to improve, and may someday eliminate the need for human transcription entirely.

Implications

The primary implications of the topic model and other associated machine learning approaches will be in (a) targeted evaluation of questions in clinical trials that compare specific therapies, and (b) exploration of large-scale naturalistic datasets that capture variability in psychotherapy as actually practiced.

First, consider a recent large ($n = 495$) clinical trial comparing psychodynamic psychotherapy with CBT for social anxiety disorder (Leichsenring et al., 2013). Both treatments were better than wait-list. Between-treatment comparisons were generally equivocal (e.g., CBT had somewhat larger remission rates, but response

rates were not significantly different; no differences met clinically significant benchmarks set a priori). Differences between therapists (5–7% of variance in outcomes) were larger than treatment effects (1–3% of variance in outcomes). As is typical with large-scale psychotherapy clinical trials, there have already been published comments (Clark, 2013) and rejoinders (Leichsenring & Salzer, 2013) on possible explanations for the findings wherein Clark raised questions about the implementation of the CBT and Leichsenring reported that the competence of psychodynamic therapists may not have been ideal. In addition, Leichsenring and Salzer (2013) noted that CBT therapists used more dynamic interventions than dynamic therapists used CBT-related interventions, raising questions about the internal validity of the trial. It is also possible that specific types of statements not specific to either intervention were responsible for between-therapist differences in outcomes.

As with other large psychotherapy clinical trials (e.g., Elkin, 1989), the debate will likely continue. However, a fundamental problem remains. While all treatment sessions were recorded, comparisons of adherence and competence were based on a total of 50 sessions (Leichsenring & Salzer, 2013). As the mean number of sessions for a patient was 25, and 416 patients received either CBT or psychodynamic treatment, the trial consisted of >10,000 sessions (seven times more sessions than included in this article). Analyses of what actually happened in this trial are driven by half of 1% of all available sessions. This sample size is typical and understandable given the labor intensiveness of behavioral coding. However, given the centrality of treatment mechanism questions to the field of psychotherapy, we look forward to more thorough analyses of process questions with computational methods. For example, researchers could conduct original human coding of subsets of sessions and use these data to train topic models that might examine a larger collection of sessions. This research may ultimately lead to more definitive answers regarding what actually happens during patient–therapist interactions and what specific therapist behaviors predict treatment outcomes within and across specific treatments.

Funding agencies may consider requiring archives of audio and transcripts for sessions in clinical trials such that they can be used in later research. Although there are privacy concerns that would need to be addressed in such a procedure, there is simply no other way for researchers to adequately evaluate what happened in the treatment. Although manuals exist, these prescriptive books are not sufficient to capture the complexity of what happens during the clinical encounter. To truly understand the mechanisms of psychotherapy, we must begin to contend with the sheer complexity and volume of linguistic data that is created during our work.

More practically, topic models could be used as adjuncts to training and fidelity monitoring in clinical trials or naturalistic settings, automatically highlighting outlier sessions or noting particular therapist interventions that were inconsistent with the specified treatment approach. In naturalistic settings, topic models could be used as a quantitatively derived aid to the traditional qualitative, report-based models of supervision. In combination with speech recognition, and selective human coding, one could imagine extremely large psychotherapy process studies (e.g., 100,000 sessions) that avoid confidentiality concerns by evaluating session content without requiring humans to listen directly to all

sessions. Studies of this size could be positioned to discover specific processes that are involved in successful versus unsuccessful cases.

Conclusions

We design treatments, package them in books, and hope that trained providers implement them in a way that is faithful to the theory and makes sense for a given patient. This implementation often involves many hours of emotional, unstructured dialogue. Specifically, the patient–provider interaction contains much of the treatment’s active ingredients. The conversation is not simply a means of developing rapport and conducting an assessment to yield a diagnosis—it *is* the treatment. As a result, the questions of interest to psychotherapy researchers are complex and imbedded in extremely large speech corpora. Research questions may include understanding the unfolding of intricate psychoanalytic concepts over a large number of sessions, the cultivation of accurate empathy, or the competent use of cognitive restructuring to examine an accurately identified irrational thought. Moreover, there is continued hope that a grand rapprochement may be possible wherein more general theories of psychotherapy process can replace and improve on the traditional encampments that have characterized the scope of psychotherapy research for two generations.

Despite the fundamentally linguistic nature of these questions, most of the raw data in psychotherapy are never subjected to empirical scrutiny. The bulk of psychotherapy process research uses patient self-report or observer ratings of provider behavior. These methods have been available for decades and have yielded important insights about the nature of psychotherapy. However, existing methods are simply not sufficient to analyze data of this size and complexity, limiting both the nuance and scale of questions that psychotherapy researchers can address. There remains an almost lawful tension between the scope and the richness of our research. One can do a very large psychotherapy study, but the data will be restricted to utilization counts and self-report measures of treatment process and clinical outcomes.

Alternatively, one can do detailed behavioral coding of sessions to evaluate therapist adherence, or qualitative work to extract themes, but the size of these studies is necessarily limited because of labor intensiveness of the work. Machine learning procedures such as the topic models used in the current study offer an opportunity to strike a balance between these poles, extracting complex information (e.g., discussions of the therapeutic relationship) on a large scale.

Most thinking about how technology will revolutionize psychotherapy focuses on the digitization of treatment itself (i.e., computer-based treatments, mobile apps, see Silverman, 2013). Many worry about how the “low tech” field of psychotherapy will adjust to this world, while more optimistic commentaries expect the technological mediation of human interaction will simply provide more grist for the mill—albeit in a different form (Tao, 2014). However, we are poised for parallel technological revolution in psychotherapy where advanced computational methods like the machine learning approach described in this article may ultimately support, query, and expand the complex, messy beauty of a therapist and patient talking.

References

- Acerbi, A., Lamos, V., Garnett, P., & Bentley, R. A. (2013). The expression of emotions in 20th century books. *PLoS ONE*, *8*, e59030. doi:10.1371/journal.pone.0059030
- Anderson, T., Bein, E., Pinnell, B., & Strupp, H. (1999). Linguistic analysis of affective speech in psychotherapy: A case grammar approach. *Psychotherapy Research*, *9*, 88–99.
- American Psychiatric Association. (2006). *American Psychiatric Association practice guidelines for the treatment of psychiatric disorders*. Arlington, VA: American Psychiatric Publishing.
- Atkins, D. C., Rubin, T. N., Steyvers, M., Doeden, M. A., Baucom, B. R., & Christensen, A. (2012). Topic models: A novel method for modeling couple and family text data. *Journal of Family Psychology*, *26*, 816–827. doi:10.1037/a0029607
- Atkins, D. C., Steyvers, M., Imel, Z. E., & Smyth, P. (in press). Automatic evaluation of psychotherapy language with quantitative linguistic models: An initial application to motivational interviewing. *Implementation Science*.
- Baardseth, T. P., Goldberg, S. B., Pace, B. T., Wislocki, A. P., Frost, N. D., Siddiqui, J. R., . . . Wampold, B. E. (2013). Cognitive-behavioral therapy versus other therapies: Redux. *Clinical Psychology Review*, *33*, 395–405. doi:10.1016/j.cpr.2013.01.004
- Benish, S. G., Imel, Z. E., & Wampold, B. E. (2008). The relative efficacy of bona fide psychotherapies for treating post-traumatic stress disorder: A meta-analysis of direct comparisons. *Clinical Psychology Review*, *28*, 746–758. doi:10.1016/j.cpr.2007.10.005
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research*, *3*, 993–1022.
- Breiman, L. (2001). Random forests. *Machine Learning*, *45*, 5–32. doi:10.1023/A:1010933404324
- Clark, D. M. (2013). Psychodynamic therapy or cognitive therapy for social anxiety disorder. *The American Journal of Psychiatry*, *170*, 1365. doi:10.1176/appi.ajp.2013.13060744
- Cox, T. F., & Cox, M. A. A. (2000). *Multidimensional scaling* (2nd ed.). Boca Raton, FL: CRC Press.
- Crits-Christoph, P., Gibbons, M. B. C., & Mukherjee, D. (2013). Psychotherapy process-outcome research. In M. J. Lambert (Ed.), *Bergin and Garfield's handbook of psychotherapy and behavior change* (pp. 298–340, 5th ed.). New Jersey: Wiley.
- Elkin, I. (1989). National Institute of Mental Health Treatment of Depression Collaborative Research Program. *Archives of General Psychiatry*, *46*, 971. doi:10.1001/archpsyc.1989.01810110013002
- Elliott, R., Bohart, A. C., Watson, J. C., & Greenberg, L. S. (2011). Empathy. *Psychotherapy*, *48*, 43–49. doi:10.1037/a0022187
- Greenberg, L. S., & Newman, F. L. (1996). An approach to psychotherapy change process research: Introduction to the special section. *Journal of Consulting and Clinical Psychology*, *64*, 435. doi:10.1037/0022-006X.64.3.435
- Griffiths, T. L., & Steyvers, M. (2004). Finding scientific topics. *Proceedings of the National Academy of Sciences United States of America*, *101* (Suppl 1), 5228. doi:10.1073/pnas.0307752101
- Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, *114*, 211–244. doi:10.1037/0033-295X.114.2.211
- Hilbert, M., & Lopez, P. (2011). The world's technological capacity to store, communicate, and compute information. *Science*, *332*, 60–65. doi:10.1126/science.1200970
- Horvath, A. O., Del Re, A. C., Flückiger, C., & Symonds, D. (2011). Alliance in individual psychotherapy. *Psychotherapy*, *48*, 9–16. doi:10.1037/a0022186
- Imel, Z. E., Barco, J. R., Brown, H. J., Baucom, B. R., Baer, J. S., Kircher, J. C., & Atkins, D. C. (2014). The association of therapist empathy and synchrony in vocally encoded arousal. *Journal of Counseling Psychology*, *61*, 146–153. doi:10.1037/a0034943
- Imel, Z. E., Wampold, B. E., Miller, S. D., & Fleming, R. R. (2008). Distinctions without a difference: Direct comparisons of psychotherapies for alcohol use disorders. *Psychology of Addictive Behaviors*, *22*, 533–543. doi:10.1037/a0013171
- Kao, J., & Jurafsky, D. (2012). A computational analysis of style, affect, and imagery in contemporary poetry. *Naacl-Hlt*, *2012*, 1–10.
- Kaplan, D. M., & Blei, D. M. (2007). A computational approach to style in American poetry. *IEEE International Conference on Data Mining*, *7*, 545–550. doi:10.1109/ICDM.2007.76
- Kirschenbaum, H. (2004). Carl Rogers's life and work: An assessment on the 100th anniversary of his birth. *Journal of Counseling and Development*, *82*, 116–124. doi:10.1002/j.1556-6678.2004.tb00293.x
- Krupski, A., Joesch, J. M., Dunn, C., Donovan, D., Bumgardner, K., Lord, S. P., . . . Roy-Byrne, P. (2012). Testing the effects of brief intervention in primary care for problem drug use in a randomized controlled trial: Rationale, design, and methods. *Addiction Science & Clinical Practice*, *7*(1), 27. doi:10.1186/1940-0640-7-27
- Lambert, M. J. (2013). *Bergin and Garfield's handbook of psychotherapy and behavior change*. New Jersey: Wiley.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240. doi:10.1037/0033-295X.104.2.211
- Lee, C. M., Kilmer, J. R., Neighbors, C., Atkins, D. C., Zheng, C., Walker, D. D., & Larimer, M. E. (2014). Indicated prevention for college student marijuana use: A randomized controlled trial. *Journal of Consulting and Clinical Psychology*, *81*, 702–709.
- Leichsenring, F., & Salzer, S. (2013). Response to Clark. *The American Journal of Psychiatry*, *170*, 1365. doi:10.1176/appi.ajp.2013.13060744r
- Leichsenring, F., Salzer, S., Beutel, M. E., Herpertz, S., Hiller, W., Hoyer, J., . . . Leibing, E. (2013). Psychodynamic therapy and cognitive-behavioral therapy in social anxiety disorder: A multicenter randomized controlled trial. *The American Journal of Psychiatry*, *170*, 759–767. doi:10.1176/appi.ajp.2013.12081125
- Manning, C., & Schütze, H. (1999). *Foundations of statistical natural language processing*. Cambridge, MA: MIT Press.
- Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*, *46*, 806–834. doi:10.1037/0022-006X.46.4.806
- Mergenthaler, E. (2008). Resonating minds: A school-independent theoretical conception and its application to psychotherapeutic processes. *Psychotherapy Research*, *18*, 109–126. doi:10.1080/10503300701883741
- Miller, W. R., & Rollnick, S. (2002). *Motivational interviewing*. New York, NY: Guilford Press.
- Morgenstern, J., & McKay, J. R. (2007). Rethinking the paradigms that inform behavioral treatment research for substance use disorders. *Addiction*, *102*, 1377–1389. doi:10.1111/j.1360-0443.2007.01882.x
- Moyers, T. B., Miller, W. R., & Hendrickson, S. M. L. (2005). How does motivational interviewing work? Therapist interpersonal skill predicts client involvement within motivational interviewing sessions. *Journal of Consulting and Clinical Psychology*, *73*, 590–598. doi:10.1037/0022-006X.73.4.590
- Neighbors, C., Lee, C. M., Atkins, D. C., Lewis, M. A., Kaysen, D., Mittmann, A., . . . Larimer, M. E. (2012). A randomized controlled trial of event-specific prevention strategies for reducing problematic drinking associated with 21st birthday celebrations. *Journal of Consulting and Clinical Psychology*, *80*, 850–862. doi:10.1037/a0029480
- Reynes, R., Martindale, C., & Dahl, H. (1984). Lexical differences between working and resistance sessions in psychoanalysis. *Journal of Clinical Psychology*, *40*, 733–737. doi:10.1002/1097-4679(198405)40:3<733::AID-JCLP2270400315>3.0.CO;2-G

- Rosen-Zvi, M., Chemudugunta, C., Griffiths, T., Smyth, P., & Steyvers, M. (2010). Learning author-topic models from text corpora. *ACM Transactions on Information Systems*, 28, 1–38. doi:10.1145/1658377.1658381
- Rubin, T. N., Chambers, A., Smyth, P., & Steyvers, M. (2012). Statistical topic models for multi-label document classification. *Machine Learning*, 88, 157–208. doi:10.1007/s10994-011-5272-5
- Salvatore, S., Gennaro, A., Auletta, A. F., Tonti, M., & Nitti, M. (2012). Automated method of content analysis: A device for psychotherapy process research. *Psychotherapy Research*, 22, 256–273. doi:10.1080/10503307.2011.647930
- Silverman, W. H. (2013). The future of psychotherapy: One editor's perspective. *Psychotherapy*, 50, 484–489. doi:10.1037/a0030573
- Steyvers, M., & Griffiths, T. (2007). Probabilistic topic models. In T. Landauer, D. McNamara, S. Dennis, & W. Kintsch (Eds.), *Latent semantic analysis*. New Jersey: LEA.
- Stiles, W. B., Shapiro, D. A., & Firth-Cozens, J. A. (1988). Verbal response mode use in contrasting psychotherapies: A within-subjects comparison. *Journal of Consulting and Clinical Psychology*, 56, 727–733.
- Talley, E. M., Newman, D., Mimno, D., Herr, B. W., Wallach, H. M., Burns, G. A. P. C., . . . McCallum, A. (2011). Database of NIH grants using machine-learned categories and graphical clustering. *Nature*, 8, 443–444. doi:10.1038/nmeth.1619
- Tao, K. W. (2014). Too close and too far: Counseling emerging adults in a technological age. *Psychotherapy*. doi:10.1037/10.1037/a0033393
- Tollison, S., Lee, C., Neighbors, C., & Neil, T. (2008). Questions and reflections: The use of motivational interviewing microskills in a peer-led brief alcohol intervention for college students. *Behavior Therapy*, 39, 183–194.
- Toutanova, K., Klein, D., Manning, C. D., & Singer, Y. (2003). *Feature-rich part-of-speech tagging with a cyclic dependency network*. In Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology, Vol. 1 (pp. 173–180). Association for Computational Linguistics, Edmonton, Canada.
- Wampold, B. E. (2001). *The great psychotherapy debate: Models, methods, and findings*. Mahwah, NJ: Erlbaum.
- Wampold, B. E. (2007). Psychotherapy: The humanistic (and effective) treatment. *American Psychologist*, 62, 855–873.
- Webb, C. A., DeRubeis, R. J., & Barber, J. P. (2010). Therapist adherence/competence and treatment outcome: A meta-analytic review. *Journal of Consulting and Clinical Psychology*, 78, 200–211.
- Williams-Baucom, K. J., Atkins, D. C., Sevier, M., Eldridge, K. A., & Christensen, A. (2010). “You” and “I” need to talk about “us”: Linguistic patterns in marital interactions. *Personal Relationships*, 17, 41–56.

Received March 2, 2014

Accepted March 3, 2014 ■