



# Automated evaluation of psychotherapy skills using speech and language technologies

Nikolaos Flemotomos<sup>1</sup> · Victor R. Martinez<sup>2</sup> · Zhuohao Chen<sup>1</sup> · Karan Singla<sup>2</sup> · Victor Ardulov<sup>2</sup> · Raghuveer Peri<sup>1</sup> · Derek D. Caperton<sup>3</sup> · James Gibson<sup>4</sup> · Michael J. Tanana<sup>5</sup> · Panayiotis Georgiou<sup>1</sup> · Jake Van Epps<sup>6</sup> · Sarah P. Lord<sup>7</sup> · Tad Hirsch<sup>8</sup> · Zac E. Imel<sup>3</sup> · David C. Atkins<sup>7</sup> · Shrikanth Narayanan<sup>1,2,4</sup>

Accepted: 15 May 2021 / Published online: 3 August 2021  
© The Psychonomic Society, Inc. 2021

## Abstract

With the growing prevalence of psychological interventions, it is vital to have measures which rate the effectiveness of psychological care to assist in training, supervision, and quality assurance of services. Traditionally, quality assessment is addressed by human raters who evaluate recorded sessions along specific dimensions, often codified through constructs relevant to the approach and domain. This is, however, a cost-prohibitive and time-consuming method that leads to poor feasibility and limited use in real-world settings. To facilitate this process, we have developed an automated competency rating tool able to process the raw recorded audio of a session, analyzing who spoke when, what they said, and how the health professional used language to provide therapy. Focusing on a use case of a specific type of psychotherapy called “motivational interviewing”, our system gives comprehensive feedback to the therapist, including information about the dynamics of the session (e.g., therapist’s vs. client’s talking time), low-level psychological language descriptors (e.g., type of questions asked), as well as other high-level behavioral constructs (e.g., the extent to which the therapist understands the clients’ perspective). We describe our platform and its performance using a dataset of more than 5000 recordings drawn from its deployment in a real-world clinical setting used to assist training of new therapists. Widespread use of automated psychotherapy rating tools may augment experts’ capabilities by providing an avenue for more effective training and skill improvement, eventually leading to more positive clinical outcomes.

**Keywords** Quality assessment · Psychotherapy · Motivational interviewing · Machine learning · Speech processing · MISC

## Need for psychotherapy quality assessment tools

Psychotherapy is a commonly used process in which mental health disorders are treated through communication between an individual and a trained mental health professional. Even though its positive effects have been well documented (Lambert & Bergin, 2002; Weisz et al., 1995; Perry et al., 1999), there is room for improvement in terms of the quality of services provided. In particular, a substantial number of patients report negative outcomes, with

signs of mental health deterioration after the end of therapy (Klatte et al., 2018; Curran et al., 2019). Apart from patient characteristics (Lambert & Bergin, 2002), therapist factors play a significant and clinically important role in contributing to negative outcomes (Saxon, Barkham, Foster, & Parry, 2017). This has direct implications for more rigorous training and supervision (Lambert & Ogles, 1997), quality improvement, and skill development. A critical factor that can lead to increased performance and thus ensure high quality of services is the provision of accurate feedback to the practitioner (Hattie & Timperley, 2007). This can take various forms; both client progress monitoring (Lambert, Whipple, & Kleinstäuber, 2018) and performance-based feedback (Schwalbe, Oh, & Zweben, 2014) have been reported to reduce therapeutic skill erosion and to contribute to improved clinical outcomes. The timing of the feedback is of utmost importance as well, since it has been shown that immediate feedback is more effective than delayed (Kulik & Kulik, 1988).

---

This research was supported by the National Institutes of Health.

✉ Nikolaos Flemotomos  
flemotom@usc.edu

Extended author information available on the last page of the article.

In psychotherapy practice, however, providing regular and immediate performance evaluation is almost impossible. Behavioral coding—the process of listening to audio recordings and/or reading session transcripts in order to observe therapists’ behaviors and skills (Bakeman & Quera, 2012)—is both time-consuming and cost-prohibitive when applied in real-world settings. It has been reported (Moyers, Martin, Manuel, Hendrickson, & Miller, 2005) that, after intensive training and supervision that lasts on average 3 months, a proficient coder would need up to 2 h to code just a 20min-long session of motivational interviewing (MI), a specific type of psychotherapy that is also the focus of the current study. The labor-intensive nature of coding means that the vast majority of psychotherapy sessions are not evaluated. As a result, many providers get inadequate feedback on their therapy skills after their initial training (Miller, Sorensen, Selzer, & Brigham, 2006) and behavioral coding is mainly applied for research purposes with limited outreach to community settings (Proctor et al., 2011). At the same time, the barriers imposed by manual coding usually lead to research studies with relatively small sample sizes (Magill et al., 2014), limiting progress in the field. It is, thus, made apparent that being able to evaluate a therapy session and provide feedback to the practitioner at a low cost and in a timely manner would both boost psychotherapy research and scale up quality assessment to real-world use. In the current work, we investigate whether it is feasible to analyze a therapy session recording in a fully automatic way and provide rich feedback to the therapist within short time.

## Behavioral coding for motivational interviewing

Motivational interviewing (MI; Miller & Rollnick 2012), often used for treating addiction and other conditions, is a client-centered intervention that aims to help clients make behavioral changes through resolution of ambivalence. It is a psychotherapy treatment with evidence supporting that specific skills are correlated with the clinical outcome (Gaume, Gmel, Faouzi, & Daeppen, 2009; Magill et al. 2014) and also that those skills cannot be maintained without ongoing feedback (Schwalbe et al. 2014). Thus, great effort from MI researchers has been devoted to developing instruments to evaluate fidelity to MI techniques.

The gold standard for monitoring clinician fidelity to treatment is behavioral observation and coding (Bakeman & Quera, 2012). During that process, trained coders assign specific labels or numeric values to the psychotherapy session, which are expected to provide important therapy-related details (e.g., “how many open questions were posed by the therapist?” or “did the counselor accept and

respect the client’s ideas?”) and essentially reflect particular therapeutic skills. While there are a variety of coding schemes (Madson & Campbell, 2006), in this study we focus on a widely used research tool, the Motivational Interviewing Skill Code (MISC 2.5; Houck, Moyers, Miller, Glynn, & Hallgren, 2010), which was specifically developed for use with recorded MI sessions (Madson & Campbell, 2006). MISC defines behavior codes both for the counselor and the patient, but for the automated system reported in this paper we focus on counselor behaviors.

The MISC manual (Houck et al. 2010) defines both session-level and utterance-level codes. The session-level (or “global”) codes characterize the entire interaction and are scored on a five-point Likert scale ranging from 1 (*poor*) to 5 (*excellent*). Table 1 gives an overview of the six therapist-related global MISC ratings with a short description for each one. When coding at the utterance-level, instead of assigning numerical values, the coder decides in which behavior category each utterance belongs. An utterance is a “thought unit” (Houck et al. 2010), which means that multiple consecutive phrases might be parsed into a single utterance and, likewise, multiple utterances might compose a single sentence or talk turn. After the session is parsed into utterances, each one is assigned one of the codes summarized in Table 2 (or gets the label NC if it can not be coded).

The platform we present is evaluated under real-world conditions, by continuously gathering and analyzing psychotherapy sessions recorded in the counseling center of an American university with a large student body. Our system is part of a broader study where the goal is to investigate whether therapists make more extensive use of MI techniques after MI-related training and we thus evaluate all the recorded sessions following the MISC protocol.

**Table 1** Therapist-related session-level codes, as defined by MISC 2.5

| Name             | High score means that counselor...  |
|------------------|---|
| acceptance       | consistently communicates acceptance and respect to the client                |
| empathy          | makes an effort to accurately understand the client’s perspective             |
| direction        | is focused on a specific target behavior                                      |
| autonomy support | does not attempt to control the client’s behavior or choices                  |
| collaboration    | interacts with their clients as partners and avoids an authoritarian attitude |
| evocation        | tries to “draw out” client’s own desire for changing                          |

Each code is scored on a five-point Likert scale

**Table 2** Therapist-related utterance-level codes, as defined by MISC 2.5

| Abbr. | Name                          | Example   |
|-------|-------------------------------|---|
| ADP   | Advise with Permission        | Would it be all right if I suggested something?                     |
| ADW   | Advise w/o Permission         | I recommend that you attend 90 meetings in 90 days.                 |
| AF    | Affirm                        | Thank you for coming today.   |
| CO    | Confront                      | (C: I don't feel like I can do this.) Sure you can.                 |
| DI    | Direct                        | Get out there and find a job.                                       |
| EC    | Emphasize Control             | It is totally up to you whether you quit or cut down.               |
| FA    | Facilitate                    | Uh huh. ( <i>keep-going acknowledgment</i> )                        |
| FI    | Filler                        | Nice weather today!   |
| GI    | Giving Information            | Your blood pressure was elevated [...] this morning.                |
| QUO   | Open Question                 | Tell me about your family.  |
| QUC   | Closed Question               | How often did you go to that bar?                                   |
| RCP   | Raise Concern with Permission | Could I tell you what concerns me about your plan?                  |
| RCW   | Raise Concern w/o Permission  | That doesn't seem like the safest plan.                             |
| RES   | Simple Reflection             | (C: The court sent me here.) That's why you're here.                |
| REC   | Complex Reflection            | (C: The court sent me here.) This wasn't your choice to be here.    |
| RF    | Reframe                       | (C: [...] something else comes up [...]) You have clear priorities. |
| SU    | Support                       | I'm sorry you feel this way.  |
| ST    | Structure                     | Now I'd like to switch gears and talk about exercise.               |
| WA    | Warn                          | Not showing up for court will send you back to jail.                |
| NC    | No Code                       | You know, I... ( <i>meaning is not clear</i> )                      |

Most of the examples are drawn from the MISC manual (Houck et al. 2010). Many of the code assignments depend on the client's previous utterance (C)

## Psychotherapy evaluation in the digital era

Psychotherapy sessions are interventions primarily based on spoken language, which means that the information capturing the session quality is encoded in the speech signal and the language patterns of the interaction. Thus, with the rapid technological advancements in the fields of Speech and Natural Language Processing (NLP) over the last few years (e.g., Devlin, Chang, Lee, & Toutanova, 2017, Devlin et al. 2019), and despite many open challenges specific to the healthcare domain (Quiroz et al. 2019), it is not surprising to see trends in applying computational techniques to automatically analyze and evaluate psychotherapy sessions.

Such efforts span a wide range of psychotherapeutic approaches including couples therapy (Black et al. 2013), MI (Xiao et al. 2016) and cognitive behavioral therapy (Flemotomos et al. 2018), used to treat a variety of conditions such as addiction (Xiao et al. 2016) and post-traumatic stress disorder (Shiner et al., 2012). Both text-based (Imel, Steyvers, & Atkins, 2015; Xiao, Can, Georgiou, Atkins, & Narayanan, 2012) and audio-based (Black et al. 2013; Xiao et al., 2014) behavioral descriptors have been explored in the literature and have been used either unimodally or in combination with each other (Singla et al., 2018).

In this study, we focus on behavior code prediction from textual data. Most research studies focused on text-based behavioral coding have relied on written text excerpts (Rojas-Barahona et al. 2018) or used manually derived transcriptions of the therapy session (Lee et al., 2019; Can et al., 2015; Gibson et al., 2019). However, a fully automated evaluation system for deployment in real-world settings requires a speech processing pipeline that can analyze the audio recording and provide a reliable speaker-segmented transcript of what was spoken by whom. This is a necessary condition before such an approach is introduced into clinical settings since, otherwise, it may eliminate the burden of manual behavioral coding, but it introduces the burden of manual transcription. Transcription errors introduced by Automatic Speech Recognition (ASR) algorithms may have a significant effect on the performance of NLP-based models (Malik et al., 2018), so demonstrating the practical feasibility of a fully automated pipeline is an important task.

An end-to-end system is presented by Xiao et al. (2015) and Xiao et al. (2016), where the authors report a case study of automatically predicting the empathy expressed by the provider. A similar platform, focused on couples therapy, is presented by Georgiou et al. (2011). Even employing an ASR module with relatively high error rate,

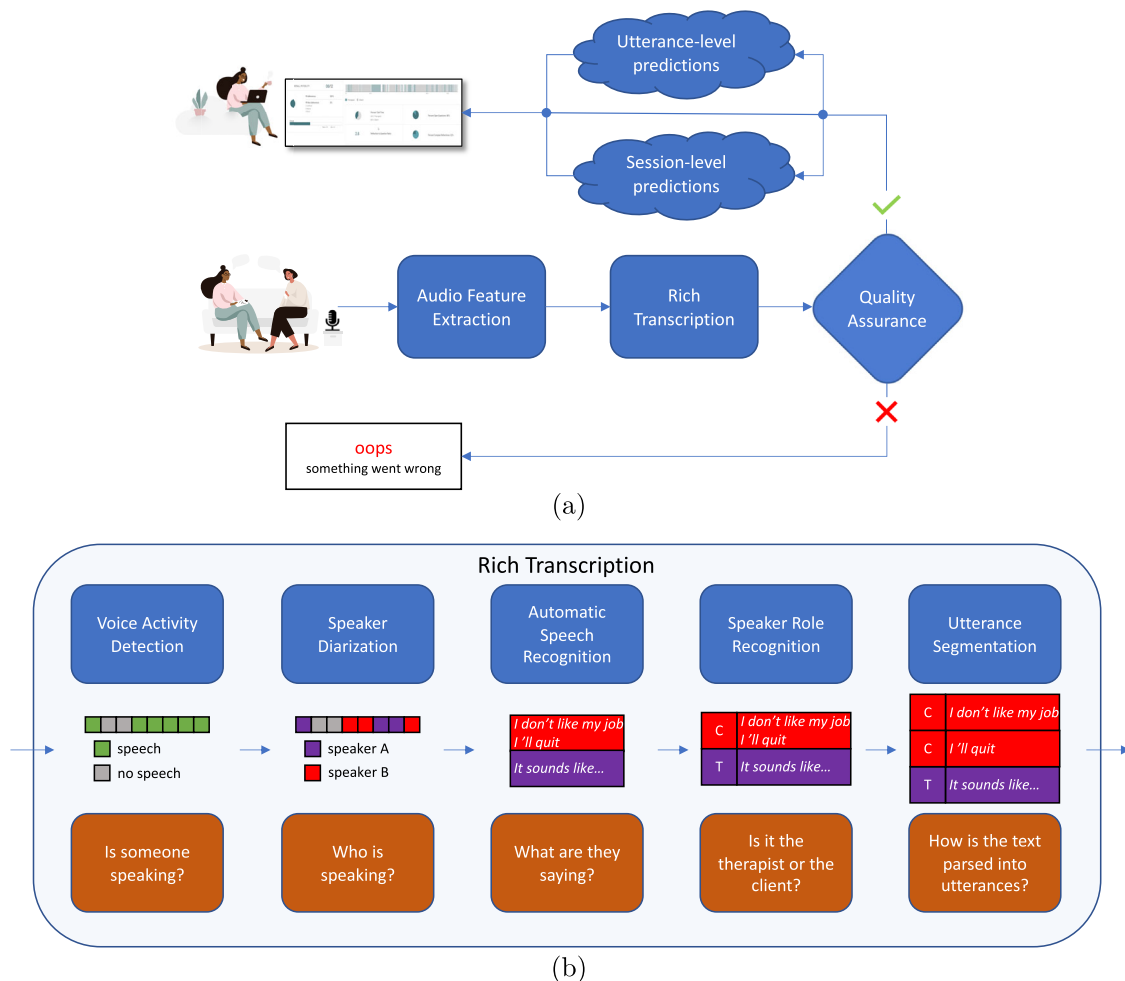
those systems were reported to provide competitive prediction performance (Georgiou et al., 2011). The scope of the particular studies, though, was limited only to session-level codes, while the evaluation sessions were selected from the two extremes of the coding scale. Thus, for each code, the problem was formulated as a binary classification task trying to identify therapy sessions where a particular code (or its absence) is represented more prominently (e.g., identify ‘low’ vs. ‘high’ empathy).

## Current study

In the current work, we demonstrate and analyze a platform (Fig. 1) able to process the raw recording of a psychotherapy

session and provide, within short time, performance-based feedback according to therapeutic skills and behaviors expressed both at the utterance and at the session level. We focus on dyadic psychotherapy interactions (i.e., one therapist and one client) and the quality assessment is based on the counselor-related codes of the MISC protocol (Houck et al., 2010). The behavioral codes are predicted by NLP algorithms that analyze the linguistic information captured in the automatically derived transcriptions of the session.

The overall architecture is illustrated in Fig. 1a. After both parties have formally consented, the therapist begins recording the session. The digital recording is directly sent to the processing pipeline and appropriate acoustic features are extracted from the raw speech signal. The rich audio transcription component of the system (Fig. 1b) consists of



**Fig. 1** **a** Overview of the system used to assess the quality of a psychotherapy session and provide feedback to the therapist. Once the audio is recorded, it is automatically transcribed to find who spoke when and what they said. If the transcription meets certain quality criteria, this textual information is used to predict utterance-level and session-level behavior codes which are summarized into an

interactive feedback report. Otherwise, an error message is displayed to the user. **b** Rich transcription module. The dyadic interaction is transcribed through a pipeline that extracts the linguistic information encoded in the speech signal and assigns each speaker turn to either the therapist or the client

five main steps: (a) Voice Activity Detection (VAD), where speech segments are detected over silence or background noise, (b) speaker diarization, where the speech segments are clustered into same-speaker groups (e.g., speaker A, speaker B of a dyad), (c) Automatic Speech Recognition (ASR), where the audio speech signal of each speaker-homogeneous segment is transcribed to words, (d) Speaker Role Recognition (SRR), where each speaker group is assigned their role: in our case study, ‘therapist’ or ‘client’, and (e) utterance segmentation, where the speaker turns are parsed into utterances which are the basic units of behavioral coding. The generated transcription is used to estimate a variety of behavior codes both at the utterance and at the session level, which reflect target constructs related to therapist behaviors and skills.

The behavioral analysis of the counselor is summarized into a comprehensive feedback report provided through an interactive web-based platform (Hirsch et al. 2018; Imel et al. 2019). Through the platform, the user is able to review the raw MISC predictions of the system (e.g., empathy score and utterances labeled as reflections), several theory-driven functionals of those (e.g., ratio of questions to reflections), session statistics (e.g., ratio of client’s to therapist’s talking time), as well as the entire speaker-segmented transcription, accompanied by the corresponding audio recording. Additionally, the user is given the option to take notes and make comments linked to specific timestamps or utterances. That way, the platform can be used directly by the provider as a self-assessment method or by a supervisor as a supportive tool that helps them deliver more effective and engaging training.

Since the system was designed with real-world deployment in mind, it was important to incorporate specific confidence metrics which reflect the quality of the automatic transcription. Employing quality safeguards helps us both identify potential computational errors, and determine whether the input was an actual therapy session or not (e.g., whether the therapist pushed the recording button by mistake). If certain quality thresholds are not met, then the final report is not generated and feedback is not provided for the specific session. Instead, an error message is displayed to the counselor. For example, in a scenario where speaker segmentation fails because the recording is too noisy or the two speakers have very similar acoustic characteristics, the system would not know which utterances correspond to the provider and which correspond to the client; as a result, the subsequent prediction algorithms would fail to accurately capture counselor-related behaviors. Being able to avoid such scenarios is of crucial importance for a system used in clinical settings.

As illustrated in Fig. 1, we have chosen a pipelined implementation of the system, as opposed to a more con-

volved architecture, potentially able to predict behavioral codes directly from the speech waveform. That way, we are able to provide a feedback report containing much richer information than merely the behavior codes or statistics of those. In particular, the user has access to the entire transcript and can understand how particular behaviors are linked to the linguistic content of the corresponding utterances. This design increases the interpretability and, as a result, the trust of the clinical provider to the system. Additionally, we are able to extract and provide information critical for the quality assessment of the therapy session, not directly related to behavior codes, such as the client’s speaking time. Finally, the quality assurance of the generated transcription is based on certain quality safeguards (described later in the paper) corresponding to specific sub-modules of the pipeline, such as the VAD and the diarization. So, if a potential error is detected at an early stage of the pipeline (e.g., VAD), the entire processing can be halted, thus avoiding wasting computational resources.

## Materials and methods

### Datasets

The design of the system presented in this work is based on datasets drawn from a variety of sources. We have combined large speech and language corpora both from the psychotherapy domain and from other fields (meetings, telephone conversations, etc.). That way, we wanted to ensure high in-domain accuracy when analyzing psychotherapy data, but also robustness across various recording conditions. In order to use and evaluate the system in real-world clinical settings, we have additionally collected and analyzed a set of more than 5000 recordings of therapy sessions between a provider and a patient at a University Counseling Center (UCC). The details of the various datasets are presented in the following sections.

### Out-of-domain corpora

**Audio sources** The acoustic modeling performed in this work was mainly based on a large collection of speech corpora, widely used by the research community for a variety of speech processing tasks. Specifically, we used the Fisher English (Cieri et al., 2004), ICSI Meeting Speech (Janin et al. 2003), WSJ (Paul & Baker, 1992), and 1997 HUB4 (Graff et al., 1997) corpora, available through the Linguistic Data Consortium (LDC), as well as Librispeech (Panayotov et al., 2015), TED-LIUM (Rousseau et al., 2014), and AMI (Carletta et al. 2005). This combined speech dataset consists of more than 2000 hours of



audio and contains recordings from a variety of scenarios, including business meetings, broadcast news, telephone conversations, and audiobooks/articles.

**Text sources** The aforementioned datasets are accompanied by manually derived transcriptions which can be used for language modeling tasks. In our case, since we need to capture linguistic patterns specific to the psychotherapy domain, the main reason we need some out-of-domain text corpus is to build a background model that guarantees a large enough vocabulary and minimizes the unseen words during evaluation. To that end, we use the transcriptions of the Fisher English corpus, featuring a vocabulary of 58.6 K words and totaling more than 21 M tokens.

### Psychotherapy-related corpora

**Audio sources** In order to train and adapt our machine learning models, used both for the transcription component of the system and for the behavior coding predictions, we also used several psychotherapy-focused corpora. In particular, we used a collection of 337 MI sessions (for which audio, transcription and manual coding information were available) from six independent clinical trials (ARC, ESPSB, ESP21, iCHAMP, HMCBI, CTT). In more detail, ARC (Tollison et al. 2008; nine sessions), ESPSB (Lee et al. 2014; 38 sessions) and ESP21 (Neighbors et al. 2012; 19 sessions) feature brief alcohol interventions. CTT (Baer et al. 2009; 19 sessions) also consists of alcohol interventions, but using standardized patients (i.e., actors portraying patients). Finally, iCHAMP (Lee et al. 2013; seven sessions) addresses marijuana addiction and HMCBI (Krupski et al. 2012; 70 sessions) addresses poly-drug abuse. We refer to the combined dataset as the TOPICS-CTT corpus and we have split it into train (TOPICS-CTT<sub>train</sub>; 242 sessions) and test (TOPICS-CTT<sub>test</sub>; 95 sessions) sets.

The mean duration of the sessions is 29.10min (std=15.65min). The number of unique therapists and clients recorded in those sessions is given in Table 3. Unfortunately, the client IDs are not available for the HMCBI sessions, so the exact total number of different clients is not known. However, under the assumption that it is highly

improbable for the same client to visit different therapists in the same study, and having the necessary metadata available for the rest of the corpus, we make the train/test split in a way that we are highly confident there is no overlap between speakers. This is important since we want to make sure that our models capture universal behavior-specific patterns during training and not speaker-specific linguistic information.

**Text sources** The transcripts of the aforementioned MI sessions were enhanced by data provided by the Counseling and Psychotherapy Transcripts Series (CPTS), available from the Alexander Street Press ([alexanderstreet.com](http://alexanderstreet.com)) via library subscription. This included transcripts from a variety of therapy interventions totaling about 300 K utterances and 6.5 M words. For this corpus, no audio or behavioral coding are available, and the data were hence used only for language-based modeling tasks.

### University counseling center data collection

Through a collaboration with the university-based counseling center of a large western university, we gathered a corpus of real-world psychotherapy sessions to evaluate the proposed system. Therapy treatment was provided by a combination of licensed staff as well as trainees pursuing clinical degrees. Topics discussed span a wide range of concerns common among students, including depression, anxiety, substance use, and relationship concerns. All the participants (both patients and therapists) had formally consented to their sessions being recorded. Study procedures were approved by the institutional review board of the University of Utah. Each session was recorded by two microphones hung from the ceiling of the clinic offices, one omni-directional and one directed to where the therapist generally sits.

Data reported in this article were collected between September 2017 and March 2020, for a total of 5097 recordings. Every time a session is recorded, it is automatically sent to the audio processing pipeline, and a performance-based feedback report is generated. We note that some of those recordings were not actually valid

**Table 3** Number of sessions, unique therapists, and unique clients in the six clinical trials composing the TOPICS-CTT corpus

|             | ARC | ESPSB | ESP21 | iCHAMP | HMCBI | CTT |
|-------------|-----|-------|-------|--------|-------|-----|
| #sessions   | 9   | 38    | 19    | 7      | 70    | 194 |
| #therapists | 3   | 15    | 8     | 5      | 15    | 132 |
| #clients    | 9   | 38    | 19    | 7      | –     | 4   |

The client IDs are not known for the HMCBI data

therapy sessions (e.g., the therapist pushed the recording button by mistake); however we do have relevant safeguards for such cases, as described later in the article. Eventually, 4268 sessions were successfully processed with a mean duration of 49.77min (std=11.50min), giving a therapy corpus totaling more than 2.8 M utterances and 28 M words (according to the automatically generated output), including sessions from at least 59 therapists and 1040 clients (there are a few sessions for which such metadata are not available).

In order to adapt and evaluate the pipeline, 188 sessions were selected to be manually transcribed and coded. The coding took place in two independent trials (one in mid-2018 and one in late 2019), with some differences in the procedure between the two. For the first coding trial (96 sessions), the transcriptions were stripped of punctuation and coders were asked to parse the session into utterances. During the second trial (92 sessions), the human transcriber was asked to insert punctuation, which was used to assist parsing. Additionally, for the second batch of transcriptions, stacked behavioral codes (more than one code per utterance) were allowed in case one of the codes is open or closed question (QUO or QUC). Because of those differences in the coding approach, we are reporting results independently for the two trials; in particular, we have split the first trial into train ( $UCC_{train}$ ; 50 sessions), development ( $UCC_{dev}$ ; 26 sessions), and test ( $UCC_{test1}$ ; 20 sessions) sets, while we refer to the second trial as the  $UCC_{test2}$  set and we only use it for evaluation. That way, we are able to monitor the robustness of the system through time, without continuously adapting to new data. For similar reasons as in the case of the TOPICS-CTT corpus, the split for the first trial was done in a way so that there is no speaker overlap between the different sets.

Each of the 188 sessions was coded by at least one of three coders. Among those, 14 sessions (from the first trial) were coded by two or three coders, so that we can have a measure of inter-rater reliability (IRR). To that end, we estimated Krippendorff's alpha (Krippendorff, 2018) for each code, a statistic that is generalizable to different types of variables and flexible with missing observations (Hallgren, 2012). Since sessions were parsed into utterances from the human raters, the unit of coding is not fixed, so we got an estimate for the utterance-level codes at the session level by using the per session occurrences of each label. For the IRR analysis, we treated the occurrences of the utterance-level codes as ratio variables and the values of the session-level codes as ordinal variables. The results for all the codes are given in Table 4. For the session-level codes, the 'within one' reliability is also provided, since it

**Table 4** Krippendorff's alpha ( $\alpha$ ) to estimate inter-rater reliability (IRR) for the utterance-level (*upper four tables*; ratio measurement level) and the session-level (*lower table*; ordinal measurement level) codes in the University Counseling Center (UCC) data

| Code          | IRR ( $\alpha$ ) | IRR 'within 1' ( $\alpha$ ) |
|---------------|------------------|-----------------------------|
| acceptance    | 0.468            | 0.747                       |
| empathy       | 0.532            | 1.000                       |
| direction     | 0.593            | 0.795                       |
| aut. support  | 0.464            | 0.743                       |
| collaboration | 0.287            | 0.472                       |
| evocation     | 0.410            | 0.626                       |
| ADP           | 0.542*           |                             |
| ADW           | 0.422            |                             |
| AF            | 0.123            |                             |
| CO            | 0.497*           |                             |
| DI            | 0.590            |                             |
| EC            | 0.558            |                             |
| FA            | 0.868            |                             |
| FI            | 0.784            |                             |
| GI            | 0.861            |                             |
| QUO           | 0.945            |                             |
| QUC           | 0.897            |                             |
| RCP           | —*               |                             |
| RCW           | 0.000*           |                             |
| RES           | 0.268            |                             |
| REC           | 0.478            |                             |
| RF            | 0.093*           |                             |
| SU            | 0.345            |                             |
| ST            | 0.434            |                             |
| WA            | -0.054*          |                             |

For the utterance-level codes (abbreviations defined in Table 2), we get an estimate through their per-session occurrences. For the session-level codes, the 'within one' agreement is also provided, demonstrating whether the distance between the raters' different scores was at most one point in the Likert scale. \* denotes that the particular code was not used (count=0) by at least 2 coders for at least half of the analyzed sessions. RCP was never used by any coder

is recommended that only a distance between the raters' different scores greater than one point in the Likert scale should be considered disagreement (Schmidt et al., 2019).

## Data pre-processing

The manually transcribed UCC sessions do not contain any timing information, which means that we needed to align the provided audio with text. That way, we were able to get estimates of the "ground truth" information required

to evaluate some of the modules of our system, such as VAD and diarization. We did so by using the Gentle forced aligner ([github.com/lowerquality/gentle](https://github.com/lowerquality/gentle)), an open-source, Kaldi-based (Povey et al. 2011) tool, in order to align at the word level. However, we should note that this inevitably introduces some error to the evaluation process, since 9.4% of the words per session on average (std=3.4%) remain unaligned.

Another pre-processing step we needed to take in order to have a meaningful evaluation of the system on the UCC data is related to the behavioral labels assigned by the humans and by the platform. In particular, some of the utterance-level MISC codes are assigned very few times within a session by the human raters and the corresponding IRR is very low (Table 4); additionally, there are pairs or groups of codes with very close semantic interpretation as reflected by the examples in Table 2 (e.g., complex reflections (REC) and reframes (RF)). Thus, we clustered the codes into composite groups resulting in nine target labels. The mapping between the codes defined in the MISC manual and the target labels, as well as the occurrences of those labels in the UCC data, is given in Table 5. Comparing Tables 4 and 5, we see that IRR is substantially higher, on average, after this grouping. The facilitate code (FA) seems to dominate the data, because most of the verbal fillers (e.g., uh-huh, mm-hmm, etc.)—which are very frequent constructs in conversational speech—and single-word utterances (e.g., yeah, right, etc.) are labeled as FA.

### Audio feature extraction

For all the modules of the speech pipeline (VAD, diarization, ASR), the acoustic representation is based on the widely used Mel-frequency cepstrum coefficients (MFCCs). For the UCC data, the channels from the two recording microphones are combined through acoustic beamforming (Anguera et al., 2007), using the open-source BeamformIt tool ([github.com/xanguera/BeamformIt](https://github.com/xanguera/BeamformIt)).

### Automatic rich transcription

Before proceeding to the automatic behavioral coding, we need to transcribe the raw audio recording, in order to get information about the content, the speakers, and the utterance boundaries. This is not just a pre-processing step allowing us to apply NLP algorithms, but it also provides invaluable information which will be later incorporated in the final feedback report (e.g., talking time of each speaker). The rich transcription pipeline we propose is illustrated in Fig. 1b. In the following sections, we describe the various sub-modules of the system. Further technical details are provided in the online supplementary material that accompanies the article (Appendix A).

### Voice activity detection

The first step of the transcription pipeline is to extract the voiced segments of the input audio session. The rest of the session is considered to be silence, music, background noise, etc., and is not taken into account for the subsequent steps. To that end, a two-layer feed-forward neural network is used giving a frame-level probability. This is a pre-trained model, initially developed as part of the Robust Automatic Transcription of Speech (RATS) program (Thomas et al., 2015). The model was trained to reliably detect speech activity in highly noisy acoustic scenarios, with most of the noise types included during training being military noises like machine gun, helicopter, etc. Hence, in order to make the model better suited to our task, the original model was adapted using the UCC<sub>dev</sub> data. Optimization of the various parameters was done with respect to the unweighted average recall. The frame-level outputs are smoothed via a median filter and converted to longer speech segments which are passed to the diarization sub-system. During this process, if the silence between any two contiguous voiced segments is less than 0.5 s, the corresponding segments are merged together.

### Speaker diarization

Speaker diarization answers the question “who spoke when” and it traditionally consists of two steps. First, the speech signal is partitioned into segments where a single speaker is present. Then, those speaker-homogeneous segments are clustered into same-speaker groups. For this work, we follow the x-vector/PLDA paradigm, an approach known to achieve state-of-the-art performance for speaker recognition and diarization (Sell et al. 2018; Snyder et al., 2018). In particular, each voiced segment, as predicted by VAD, is partitioned uniformly into subsegments and for each subsegment a fixed-dimensional speaker embedding (x-vector) is extracted. Once the x-vectors have been extracted, an affinity matrix is constructed with the pairwise distances between the subsegments. The similarity metric used is based on the probabilistic linear discriminant analysis (PLDA) framework (Ioffe, 2006; Prince & Elder, 2007) within which each data point is considered to be the output of a model which incorporates both within-individual and between-individual variation. The subsegments are finally clustered together according to hierarchical agglomerative clustering (HAC). The assumption here is that each session has exactly two speakers (i.e., therapist vs. client), so we continue the HAC procedure until two clusters have been constructed. As a post-processing step after diarization, adjacent speech segments assigned to the same speaker are concatenated together into a single speaker turn, allowing a maximum of 1sec in-turn silence.



## Automatic speech recognition

After we get the speaker-homogeneous segments from the diarization module, we need to extract the linguistic content captured within each segment, since this will be the information supplied to the subsequent text-based algorithms. ASR depends on two components; the acoustic model (AM), which calculates the likelihood of acoustic observations given a sequence of words, and the language model (LM), which calculates the likelihood of a word sequence by describing the distribution of typical language usage.

In order to train the AM, we build a time-delay neural network (TDNN) with subsampling (Peddinti et al., 2015), an architecture which has been successfully applied in conversational speech achieving remarkable performance (Peddinti et al., 2015). The network is trained on a large combined speech dataset composed of the Fisher English, ICSI Meeting Speech, WSJ, 1997 HUB4, Librispeech, TED-LIUM, AMI, and TOPICS-CTT corpora. Among those, TED-LIUM and the clean portion of Librispeech are augmented with speed perturbation, noise, and reverberation (Ko et al., 2015). The final combined, augmented corpus contains more than 4000 h of phonetically rich speech data, recorded under different conditions and reflecting a variety of acoustic environments. The ASR AM was built and trained using the Kaldi speech recognition toolkit (Povey et al. 2011).

In order to build the LM, we independently train two 3-gram models using the SRILM toolkit (Stolcke, 2002). One is trained with in-domain psychotherapy data from the CPTS transcribed sessions. This is interpolated with a large background model, in order to minimize the unseen words during the system employment. The background LM is trained with the Fisher English corpus, which features conversational telephone data.

## Speaker role recognition

After diarization has been performed, we have the entire set of utterances clustered into two groups; however, there is not a natural correspondence between the cluster labels and the actual speaker roles (i.e., therapist and client). For our purposes, speaker role recognition (SRR) is exactly the task of finding the mapping between the two. Even though different speaker roles follow distinct patterns across various modalities (e.g., audio, language, structure), the linguistic stream of information is often the most useful for the task in hand (Flemotomos et al., 2018). So, in this work we are focusing on this modality, provided by the ASR output.

Let's denote the two clusters which have been identified by diarization as  $S_1$  and  $S_2$ , each one containing the utterances assigned to the two different speakers. We know a priori that one of those speakers is the therapist (T) and one is the client (C). In order to do the role matching, two trained LMs, one for the therapist ( $LM_T$ ) and one for the client ( $LM_C$ ), are used. We then estimate the perplexities of  $S_1$  and  $S_2$  with respect to the two LMs and we assign to  $S_i$  the role that yields the minimum perplexity. In case one role minimizes the perplexity for both speakers, we first assign the speaker for whom we are most confident. The confidence metric is based on the absolute distance between the two estimated perplexities (Flemotomos et al., 2018). The required LMs are 3-gram models trained with the SRILM toolkit (Stolcke, 2002), using the TOPICS-CTT<sub>train</sub> and CPTS corpora.

## Utterance segmentation

The output of the ASR and SRR modules is at the segment level, with the segments defined by the VAD and diarization

**Table 5** Mapping between the MISC-defined behavior codes (abbreviations defined in Table 2) and the grouped target labels, together with the occurrences of each group in the training and development University Counseling Center (UCC) sets

| Group | MISC codes                     | count (UCC <sub>train</sub> ) | count (UCC <sub>dev</sub> ) | IRR ( $\alpha$ ) |
|-------|--------------------------------|-------------------------------|-----------------------------|------------------|
| FA    | FA                             | 5581                          | 2500                        | 0.868            |
| GI    | GI, FI                         | 3797                          | 1695                        | 0.898            |
| QUC   | QUC                            | 1911                          | 693                         | 0.897            |
| QUO   | QUO                            | 1116                          | 405                         | 0.946            |
| REC   | REC, RF                        | 2212                          | 1041                        | 0.479            |
| RES   | RES                            | 609                           | 155                         | 0.268            |
| MIN   | ADP, ADW, CO, DI, RCW, RCP, WA | 479                           | 163                         | 0.606            |
| MIA   | AF, EC, SU                     | 428                           | 238                         | 0.363            |
| ST    | ST                             | 542                           | 135                         | 0.434            |

MIA stands for MI-adherent codes and MIN for MI-non adherent. The inter-rater reliability (IRR) for the grouped labels is also given in terms of the Krippendorff's alpha ( $\alpha$ ) value

algorithms. However, silence and speaker changes are not always the right cues to help us distinguish between utterances, which are the basic units of behavioral coding. The presence of multiple utterances per speaker turn is a challenge we often face when dealing with conversational interactions. Especially in the psychotherapy domain, it has been shown that the right utterance-level segmentation can significantly improve the performance of automatic behavior code prediction (Chen et al., 2020).

Thus, we have included an utterance segmentation module at the end of the automatic transcription, before employing the subsequent NLP algorithms. In particular, we merge together all the adjacent segments belonging to the same speaker in order to form speaker-homogeneous talk-turns, and we then segment each turn using the DeepSegment tool ([github.com/notAI-tech/deepsegment](https://github.com/notAI-tech/deepsegment)). DeepSegment has been designed to perform sentence boundary detection having specifically ASR outputs in mind, where punctuation is not readily available. In this framework, sentence segmentation is viewed as a sequence labeling problem, where each word is tagged as being either at the beginning of a sentence (utterance), or anywhere else. DeepSegment addresses the problem employing a bidirectional long-short term memory (BiLSTM) network with a conditional random field (CRF) inference layer (Ma & Hovy, 2016).

### Quality assurance

The goal of the current study is to provide accurate and reliable feedback to the counselor in a real-world environment. Hence, it is essential that we ensure we do not produce feedback reports which are problematic, either because of bad audio quality, or because of errors during computation. We have identified that most of the errors are produced during the first steps of the processing pipeline and are propagated to the subsequent steps. Thus, we have incorporated simple quality safeguards, able to catch errors associated with the audio recording, the VAD, or the diarization. Specifically, before any further processing, the following conditions need to be met:

1. The duration of the entire recording has to be between *60sec* and *5h*. Given that a typical therapy session in our study is about *50min* long, a session outside this range indicates either that the provider pushed the recording button by mistake, or that they forgot to stop recording.
2. At least 25% of the session has to be flagged as voiced, according to the VAD output. During a typical conversational interaction, there are pauses of silence which are especially useful in psychotherapy (Levitt, 2001). Although silence is an essential aspect of communication, the distribution of the silence gaps duration is highly skewed with most of them being very

short (Heldner & Edlund, 2010). If most of the therapy session is flagged as unvoiced, this is an indication of bad audio quality, of some inherent error of the VAD algorithm employed, or of a prolonged audio file where the therapist forgot to stop the recording after the actual session.

3. The average duration of the voiced segments cannot be longer than 20 s. Even though words are the primary means of communication, silence gaps are not just useful, but necessary in order for spoken language to be meaningful and natural. When our VAD system is incapable of detecting unvoiced segments, it is usually an indication of bad audio quality.
4. The minimum percentage of speech assigned to each speaker is 10% of the total speaking time. Since we are dealing with dyadic conversational scenarios, it is expected that each of the two speakers talks for a substantial amount of time. Even though therapy is not a normal dialog and the provider often plays more the role of the listener (Hill, 2009), if a person seems to be talking for less than 10% of the time (e.g., less than about *5min* in a typical *50min*-long session), then we are highly confident there is some problem. This may be an issue associated either with the audio quality, or with high speaker error introduced by the diarization module because the two speakers have similar acoustic characteristics.

If any of the aforementioned conditions is violated, the processing is halted and an error message is displayed to the end user, instead of the actual report.

### Utterance-level and session-level labeling

Once the entire session is transcribed at the utterance level, we are able to employ text-based algorithms for the task of behavior code prediction. Both utterance-level and session-level behavior codes are predicted and provided back to the counselor as part of the feedback report, as described below.

#### Utterance-level code prediction

We are focusing on counselor behaviors, so we only take into account the utterances assigned to the therapist according to the speaker role recognition. Each one of those needs to be assigned a single code from the nine target labels summarized in Table 5. This is achieved through a BiLSTM network with attention mechanism (Singla et al., 2018), which only processes textual features. The input to the system is a sequence of word-level embeddings for each utterance. The recurrent layer exploits the sequential nature of language and produces hidden vectors which take into account the entire context of each word within the utterance.

The attention layer can then learn to focus on salient words carrying valuable information for the task of code prediction, thus enhancing robustness and interpretability. The network was first trained on the TOPICS-CTT data using class weights to handle the problem of skewed code distribution in the data (Table 5). The system was further fine-tuned by continuing training on the  $UCC_{train}$  data in order to be suitably fitted to the University Counseling Center conditions.

### Session-level code prediction

Apart from the utterance-level codes, our system assigns a score to each one of the global codes of Table 1, ranging from 1 to 5. To that end, we represent the entire session, using the utterances assigned to both the therapist and the client, by the term frequency - inverse document frequency (tf-idf; Salton & McGill 1986) transformation of unigrams, bigrams, and trigrams found within the session, excluding common stop words. Those features are  $l_2$ -normalized and passed to a support vector regressor (SVR), which gives the final prediction. After hyper-parameter tuning, we chose polynomial SVR kernel (4th-degree) for acceptance and autonomy support, linear kernel for empathy, collaboration and evocation, and gaussian kernel for direction.

Contrary to the training approach followed for the utterance-level codes, here we train using only UCC data. The reason is that there is a discrepancy between the globals assigned by human raters to the TOPICS-CTT and the UCC sessions, since different coding procedures were followed. In particular, the TOPICS-CTT sessions were coded only across two global codes (empathy and MI spirit) following the Motivational Interviewing Treatment Integrity (MITI; Moyers et al. 2016) coding scheme. Thus, due to the limited amount of training data (only 188 sample points—UCC sessions—in total), we apply a five-fold cross validation scheme across the UCC dataset (from both coding trials) for any hyper-parameter tuning and we then keep those parameters to re-train using the entire UCC set.

### Final report

After we have the automatically generated transcript and all the session-level and utterance-level predictions through our system, those are provided to the therapist as a feedback report through an interactive, web-based platform which we refer to as the Counselor Observer Ratings Expert for Motivational Interviewing (CORE-MI; Hirsch et al. 2018, Imel et al. 2019). A video demonstration of the platform and its functionality is available at [www.youtube.com/watch?v=9fuvT9\\_azgw](https://www.youtube.com/watch?v=9fuvT9_azgw).

CORE-MI features two main views, the session view and the report view (Supplementary material, Appendix C,

Figure C1). In the first one, the user can listen to the recording of the therapy, watch the video (if available) and read the generated transcript, which is scrollable and searchable. Additionally, they can keep notes linked to specific timestamps and utterances of the session.

The report view provides the actual therapy session evaluation. The entire session timeline is presented in a form of a bar where talk turns of the two speakers are displayed in different colors. Hovering over a specific turn brings up the corresponding transcription and—in case the turn is assigned to the therapist—the corresponding MISC code(s). Based on the results reported later, we have decided to collapse the simple and complex reflections into one composite reflection (RE) label. The global behavior codes are also displayed, as well as a set of summary indicators which reflect the adherence to MI therapeutic skills. Those are the ratio of reflections (simple and complex) to questions (open and closed), the percentage of the open questions asked (among all the questions), the percentage of the complex reflections (among all the reflections), the percentage of each speaker's talking time, the MI adherence and the overall MI fidelity. MI adherence reflects the percentage of utterances where the counselor used MI-consistent techniques (e.g., asking open questions or giving advice with permission). Finally, the overall MI fidelity score is a composite metric rated on a 12-point scale that takes all the above into consideration and reflects the proficiency of the counselor to the different aspects of MI therapy. In particular, a provider can receive one point for passing pre-defined basic proficiency benchmarks and two points for passing advanced competency benchmarks across the following six measures of quality: empathy, MI spirit, reflection-to-question ratio, percentage of open questions, percentage of complex reflections, MI adherence. MI spirit is estimated as the average of evocation, collaboration, and autonomy support (Houck et al., 2010).

The main design characteristics of the CORE-MI platform have been tested in a past study (Hirsch et al. 2018; Imel et al. 2019) and results showed that the providers find the system easy to use and the feedback easy to understand. Additionally, most of the professional therapists that participated in the survey seemed excited about the potential opportunity to use such a system in clinical practice.

## Results and discussion

### Automatic rich transcription

All the submodules of the transcription pipeline are evaluated on the two UCC test sets we have described ( $UCC_{test1}$ ,  $UCC_{test2}$ ), both individually and as part of the overall system. That way, we want to evaluate the performance of each

one of the models, but, more importantly, investigate any error propagation that inevitably takes place.

### VAD/diarization

During evaluation, VAD is usually viewed as part of a diarization system (e.g., Sell et al. 2018), so for evaluation purposes we consider our diarization model as the first component of the pipeline (frame-level VAD results are provided in the online supplementary material - Appendix B). The standard evaluation metric for diarization is called diarization error rate (DER; Anguera et al. 2012) and it incorporates three sources of error: false alarms, missed speech, and speaker error. False alarm speech (the percentage of speech in the output but not in the ground truth) and missed speech (the percentage of speech in the ground truth but not in the output) are mostly associated with VAD. Speaker error is the percentage of speech assigned to the wrong speaker cluster after an optimal mapping between speaker clusters and true speaker labels. We estimate the DER on the UCC data using the `md-eval` tool which was developed as part of the rich transcription (RT) evaluation series ([www.nist.gov/itl/iad/mig/rich-transcription-evaluation](http://www.nist.gov/itl/iad/mig/rich-transcription-evaluation)). We have used a forgiveness collar of  $0.25\text{sec}$  around each speaker boundary, which is a standard practice (Anguera et al., 2012), and the results are reported in Table 6.

Even though the speaker confusion (speaker error rate) is on average low enough (lower than 8%), we should note that a per session analysis revealed that there are a few sessions where it is even higher than 45%. This means that diarization essentially failed for this handful of sessions, even though the human transcribers did not report any particular issues related, for example, to audio quality. Out of the three DER components, false alarm contributes most to the overall error, while the missed speech is minimal. Such a behavior is expected because of the specific implementation followed. In particular, we chose to concatenate together adjacent speech segments assigned to the same speaker, if there is not a silence gap between them greater than  $1\text{sec}$ . This step degrades the diarization result, since it labels short non-voiced segments as belonging to some speaker, thus introducing false alarms. However, it creates longer speaker-homogeneous segments, which is beneficial to ASR, and, hence, to the overall system.

### Automatic speech recognition

The evaluation of an ASR system is usually performed through the word error rate (WER) metric which is the normalized Levenshtein distance between the ASR output and the ground truth transcript. This includes errors because of word substitutions, word deletions, and word insertions.

For instance, word insertion rate is the number of new words included in the prediction which are not found in the original transcript over the total number of ground truth words. WER is calculated as the sum of those three error rates. Those errors are typically estimated for each utterance which is given to the ASR module and then summed up for all the evaluation data, in order to get an overall WER. However, when we analyze an entire therapy session which has been processed by the VAD and diarization sub-systems, the “utterances” are different than the ones identified by the human transcriber. In that case, we perform the evaluation at the session level, ignoring the speaker labels (from diarization) and concatenating all the utterances of the session. We do the same for the original transcript and we hence view the entire session as a “single utterance” for the purposes of ASR evaluation. The results are reported in Table 7.

As we can see, ASR performance is not severely degraded by any error propagation from the pre-processing step of diarization (WER is increased about 1% absolute). Interestingly, even though the insertion rate is increased, the deletion rate is decreased when the machine-generated segments are provided. This is explained by the long segments constructed by the diarization algorithm and the post-processing of its output after concatenating consecutive segments. On the one hand, labeling silence or noise as “speech” associated with some speaker occasionally leads ASR to predict words where in reality there is no speech activity—thus increasing the insertion rate. On the other hand, this minimizes the probability of missing some words because of missed speech. Such deleted words may occur when providing the oracle segments because of inaccuracies during the construction of the “ground truth” through forced alignment.

We note that, even though the estimated error is high, WERs in the range reported (30% – 40%) and even higher are typical in spontaneous medical conversations (Kodish-Wachs et al., 2018). Error analysis revealed that those numbers are inflated because of fillers (e.g. uh-huh, hmm) and other idiosyncrasies of conversational speech. It should be additionally highlighted that WER is a generic metric that gives equal importance to all the words, while for our end goal of behavior coding there are specific linguistic constructs which potentially carry more valuable information than others.

### Speaker role recognition

The described SRR algorithm operates at the session level, which means that, for evaluation purposes, it suffices to examine how many sessions are labeled correctly with respect to speaker roles. When oracle diarization information is provided, coupled either with the manual

transcriptions or with the ASR results, our algorithm achieves a perfect recognition result for all the  $UCC_{test_1}$  and  $UCC_{test_2}$  sessions. When speaker segmentation and clustering is performed through the diarization algorithm of the processing system, the SRR module fails to find the right mapping between roles and speakers for seven sessions from the  $UCC_{test_2}$  set.

This behavior is associated with error propagation from the previous steps which is made apparent from the fact that the speaker error rate for those seven sessions is 42.5% on average (std=8.5%). Given the fact that we are dealing with dyadic conversational interactions, such a high speaker confusion essentially means that the diarization algorithm failed to sufficiently distinguish between the two speakers, probably because of similar acoustic characteristics. Thus, there is not enough reliable speaker-specific linguistic information that the SRR can use during the role assignment. This example of error propagation also highlights the need for quality assurance through specific safeguards at the early steps of the processing pipeline.

### Utterance segmentation

The last step of the transcription pipeline is the utterance segmentation, which provides the basic units for behavioral coding. We get a rough indication of the quality of our segmentation process by estimating the correlation between the total number of utterances per session that have been assigned to the therapist by the human annotators and by the processing pipeline.

The Spearman correlation between them, when all the  $UCC_{test_1}$  and  $UCC_{test_2}$  sessions are taken into account, is 0.478 ( $p < 10^{-7}$ ). The number of the manually defined utterances is usually higher than the number of the ones identified by our system, because the automated rich transcription module often fails to capture very short utterances (e.g., ‘yeah’, ‘right’, etc.).

### Quality assurance

According to the quality safeguards introduced, 16 out of the 112 sessions are flagged as “problematic” in our

combined test set of  $UCC_{test_1}$  and  $UCC_{test_2}$ . All of those do not meet the fourth condition, related to the minimum allowed speaking time attributed to each speaker. This means that in practice the processing would halt after the end of the diarization algorithm, with an error message displayed to the user. When we ran the entire set of 5097 UCC recordings through the pipeline, 4268 met all the four criteria and were successfully processed.

It is interesting that, after excluding the “problematic” sessions from the test sets ( $UCC_{test_1}$ ,  $UCC_{test_2}$ ), the Spearman correlation between the total number of therapist utterances per session as assigned by the human coders vs. by the automated system is increased from 0.478 to 0.639. This is explained by the fact that, in several of those cases, poor diarization performance led the subsequent role recognition module to assign almost the entire session to the client. As a result, the number of therapist-attributed utterances was much smaller than expected.

### Utterance-level and session-level labeling

In the following sections, we discuss the results of the MISC code (utterance-level and session-level) prediction models. As in the case of the transcription pipeline submodules, we examine the effectiveness of the proposed models, both when provided with oracle information and when being part of the end-to-end system.

#### Utterance-level code prediction

When we use the manually transcribed data to perform utterance-level MISC code prediction, the overall  $F_1$  score is 0.524 for the  $UCC_{test_1}$  and 0.514 for the  $UCC_{test_2}$  sets. The  $F_1$  scores for each individual code are reported in Table 8. As expected, the results are better for the highly frequent codes (Table 5), such as the one expressing facilitation (FA), since the machine learning models have more training examples to learn from. On the other hand, the models do not perform as well for less frequent codes, such as MI-adherent and MI-non adherent behaviors (MIA and MIN). However, comparing Table 8 and Table 4, we can also see that for several of the codes that our system performs relatively poorly (e.g., simple reflections [RES]),

**Table 6** Diarization results (%) for the test sets of the University Counseling Center (UCC) data

| Set            | False alarm | Missed speech | Speaker error | DER  |
|----------------|-------------|---------------|---------------|------|
| $UCC_{test_1}$ | 13.7        | 0.4           | 6.9           | 21.0 |
| $UCC_{test_2}$ | 9.3         | 0.5           | 7.8           | 17.7 |

Diarrization error rate (DER) is estimated as the sum of false alarm, missed speech, and speaker error rates



**Table 7** Automatic speech recognition (ASR) results (%) for the test sets of the University Counseling Center (UCC) data: substitution, insertion, and deletion rates, together with the total word error rate (WER), estimated as the sum of those three

| Set                  | Diarization method | subs | del  | ins | WER  |
|----------------------|--------------------|------|------|-----|------|
| UCC <sub>test1</sub> | oracle             | 18.3 | 15.3 | 3.5 | 37.1 |
|                      | pipeline           | 20.0 | 13.9 | 4.3 | 38.1 |
| UCC <sub>test2</sub> | oracle             | 14.3 | 13.8 | 2.3 | 30.4 |
|                      | pipeline           | 16.1 | 12.5 | 3.1 | 31.6 |

Results are reported when using either the machine-generated segments (pipeline) or the ones derived by the manual transcriptions (oracle)

MI-adherent [MIA], structure [ST]), the inter-annotator agreement is also considerably low. A notable example which does not follow this pattern is the non-adherent behavior (MIN) where our system achieves the lowest results among all the codes, while there is a substantial inter-annotator agreement ( $\alpha = 0.606$ ). This is partly because of the underrepresentation of the particular code (or cluster of codes) in the training and development sets. It may also be the case that pure linguistic information found in textual patterns may not be enough for the operationalization of the particular code. This example suggests that a hybrid approach where machine learning methods are combined with knowledge-based rules from the coding manuals may be an interesting direction for future research. Finally, by examining the confusion matrices (not reported in this article), we realized that the system gets confused between the codes representing questions (QUC vs. QUO) and reflections (RES vs. REC), since those pairs of codes get usually assigned to utterances with several structural and semantic similarities.

The performance evaluation of the system when used within the pipeline is not straightforward, since the utterances given to the MISC predictor after the automatic transcription are not the same as the ones defined by the human transcribers. In that case, we use as a simple evaluation metric the correlation between the counts of each MISC label in the manual coding trial and in the automatically generated report. The results are illustrated in Fig. 2. There is a statistically significant ( $p < 0.01$ ) positive correlation for all the codes, apart from FA. The Spearman correlation for the nine codes is on average 0.446 (std = 0.136), while if we don't take into consideration the sessions that did not meet the quality criteria, the

correlation is increased to 0.566, on average (std = 0.172).

The relatively low correlation and discrepancy in the counts between the manual and the automatically generated output for FA is striking, especially if we take into account the remarkably good results of the system when we do not use the entire pipeline (Table 8). The reason is that FA is assigned to a lot of one-word utterances and talk turns. Our speech pipeline, however, often fails to capture turns of such short duration, which results in a smaller than expected frequency for the specific code. Another observed inconsistency is related to the code for simple reflections (RES), which seems to be assigned by our algorithm much more frequently than it actually occurs in the manually annotated data. As already mentioned, this is partly due to increased confusion between simple and complex reflections (RES and REC). This becomes apparent if we merge all the reflections into one composite group (denoted as RE in Fig. 2).

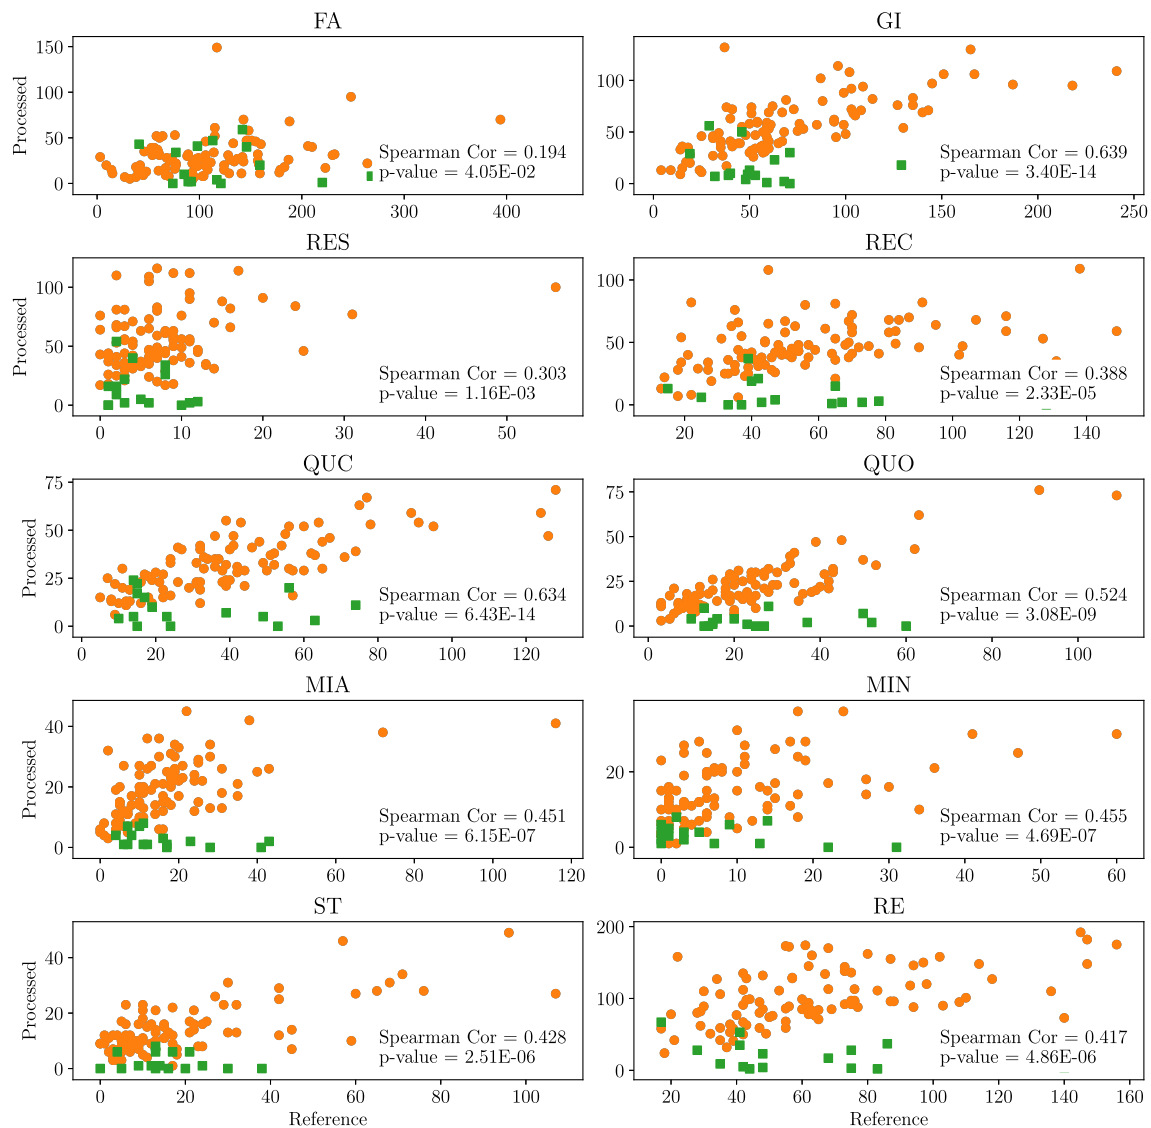
The distribution of the MISC codes across all the 4268 psychotherapy sessions that were successfully processed for this study is given in Fig. 3. As observed, the distribution is similar to the corresponding distribution if only the transcribed sessions included in the test sets (UCC<sub>test1</sub> and UCC<sub>test2</sub>) are taken into consideration. This suggests that our test sets are indicative of the entire dataset and the evaluation analysis presented likely extends to previously unseen therapy sessions processed by the system.

### Session-level code prediction

As mentioned in the Materials and methods section, the session-level code predictor is the only model where, due

**Table 8**  $F_1$  scores for the predicted utterance-level codes (Table 5) using the manually transcribed University Counseling Center (UCC) data

| Set                  | FA    | GI    | QUC   | QUO   | REC   | RES   | MIN   | MIA   | ST    |
|----------------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| UCC <sub>test1</sub> | 0.956 | 0.519 | 0.702 | 0.825 | 0.531 | 0.265 | 0.158 | 0.314 | 0.449 |
| UCC <sub>test2</sub> | 0.951 | 0.462 | 0.588 | 0.786 | 0.465 | 0.186 | 0.273 | 0.439 | 0.474 |




**Fig. 2** Count of each target MISC label per session (Table 5) when coded by humans (reference) and when processed by the pipeline. All the sessions in the two test sets of the University Counseling Center (UCC) dataset ( $UCC_{test_1}$  and  $UCC_{test_2}$ ) are shown and the

correlation values are calculated based on all of them. The sessions flagged as problematic by the quality safeguards are denoted by square markers. RE is a composite label containing both simple and complex reflections (RES and REC)

to the limited amount of training data, we apply a five-fold cross validation scheme across the entire coded UCC dataset (all 188 sessions). The cross-validation results are reported in Table 9. Results are given in terms of accuracy and averaged  $F_1$  score, after the output of SVR is rounded to the closest integer in the range from 1 to 5 and after we collapse classes 1 and 2 together (due to the very limited number of sessions scored as 1 in the reference data). We also report the ‘within one’ accuracy, demonstrating whether the distance between the reference and predicted scores was at most one. In general, the predictive power of the models seems to be lower for the codes where the inter-rater reliability (Table 4)

is also low. Additionally, the performance is not severely affected by the usage of the speech pipeline, when compared to using the manual transcriptions.

The distributions of the six global codes across all the 4268 psychotherapy sessions that were successfully processed are given in Fig. 4. All the codes, with the exception of direction, are skewed towards the higher scores of the scale (higher than 3). As was the case with the utterance-level codes (Fig. 3), we get a very similar distribution if we illustrate the results only for the sessions in the UCC test sets for which manual transcription and behavior coding information were available. This is



**Fig. 3** Frequency of the utterance-level MISC codes (Table 5) for all the University Counseling Center (UCC) recordings processed and for the subset included in the UCC test sets. Only the sessions successfully processed (that met our quality criteria) are taken into consideration

indicative of the generalization of the system and its performance to future therapy sessions.

## Limitations and conclusions

In this article, we presented and analyzed a processing pipeline able to automatically evaluate recorded psychotherapy sessions. The application of such a system in real-world settings could guarantee the provision of fast and low-cost feedback. Performance-based feedback is an essential aspect both for training new therapists and for maintaining acquired skills, and can eventually lead to improved quality of services and more positive clinical outcomes. Additionally, being able to record, transcribe, and code interventions at large scale opens up ample opportunities for psychotherapy research studies with increased statistical power.

At the point of writing, we have processed a collection of more than 5000 recordings, 4268 of which met our quality criteria and are now accompanied by transcriptions and behavioral coding information. Both utterance-level and

here. The total number of therapist-assigned utterances is about 1.2 M for all the sessions (4269 sessions) and 28 K for only the sessions included in the UCC test sets ( $UCC_{test1}$  and  $UCC_{test2}$ ; 96 sessions)

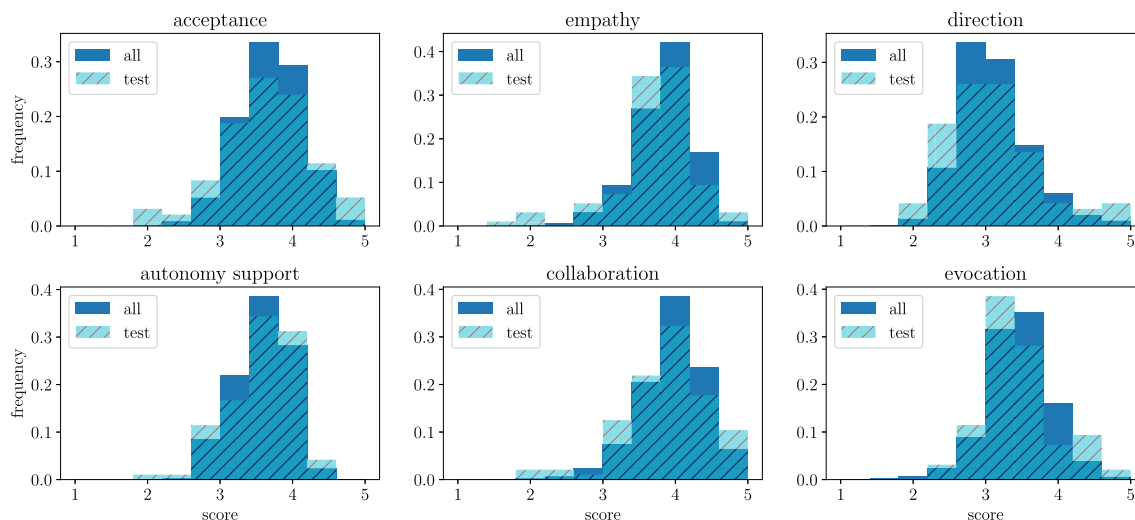
session-level MISC codes are available covering a wide range of behaviors (Figs. 3 and 4). As we are planning on expanding our corpus with more data, we are confident that such a dataset will lead to novel interesting studies in the fields of psychotherapy, computational modeling, and their intersection. For example, the transcriptions of a subset of those data have been already used to study therapeutic alliance directly using text-based features (Goldberg et al. 2020) or modeling clients and therapists as narrative characters (Martinez et al. 2019). Even though we have here focused on motivational interviewing, the basic ideas of the speech processing pipeline remain the same for other dyadic interactions as well. For instance, the same modules analyzed in this article have been used to automatically transcribe and subsequently analyze cognitive behavior therapy sessions (Chen et al., 2020).

Despite the promising results presented here, we recognize that there is room for improvement in almost all the sub-modules of the pipeline. Our analysis showed that diarization failed for some of the sessions that human transcribers had no problem processing. Additionally, there was a consistent underrepresentation of verbal fillers

**Table 9** Averaged  $F_1$  scores and accuracy for the predicted session-level MISC codes (Table 1) using the manually transcribed (oracle) or the pipeline-generated data, based on a five-fold cross validation scheme across all the University Counseling Center (UCC) test data

| Metric           | $F_1$  |          | accuracy |          | accuracy ('within 1') |          |
|------------------|--------|----------|----------|----------|-----------------------|----------|
|                  | oracle | pipeline | oracle   | pipeline | oracle                | pipeline |
| acceptance       | 0.318  | 0.297    | 0.478    | 0.457    | 0.771                 | 0.755    |
| empathy          | 0.342  | 0.342    | 0.586    | 0.580    | 0.819                 | 0.851    |
| direction        | 0.380  | 0.261    | 0.426    | 0.389    | 0.740                 | 0.697    |
| autonomy support | 0.303  | 0.261    | 0.495    | 0.451    | 0.878                 | 0.840    |
| collaboration    | 0.285  | 0.199    | 0.437    | 0.346    | 0.654                 | 0.612    |
| evocation        | 0.274  | 0.188    | 0.362    | 0.335    | 0.751                 | 0.671    |

The 'within one' accuracy demonstrates whether the distance between the predicted and reference scores was at most one point in the Likert scale



**Fig. 4** Distribution of the session-level MISC codes (Table 1) for all the University Counseling Center (UCC) recordings processed and for the subset included in the UCC test sets. Only the sessions successfully processed (that met our quality criteria) are taken into consideration here

(e.g., uh-huh) and the relevant MISC label (FA) in the automatically generated transcripts, as a result of the system struggling to capture and transcribe very short speaker turns. Moreover, the architecture design followed, where the various modules are trained independently and are then connected to form a pipeline, inevitably leads to error propagation. There are indications that alternative frameworks could reduce errors in specific cases, if for example diarization is aware of the different speaker roles (Flemotomos et al., 2020) or if the two tasks of diarization and role recognition are performed simultaneously (Flemotomos et al., 2018).

For this work, we only used text-based methods for behavioral coding. Acoustic features, however, and especially prosodic cues, play a major role in understanding language (Cutler et al., 1997) and have been successfully used in the past for MISC code prediction (Singla et al., 2018; Xiao et al., 2014). Recent studies have even shown that audio-only approaches, where word embeddings are directly learnt from spoken language, can yield improved results (Singla et al., 2020). Additionally, for the most part of our analysis, we have focused only on therapist characteristics. However, specific dialog attributes, such as speech rate entrainment (Xiao et al., 2015) and language synchrony (Lord et al., 2015; Nasir et al., 2019) between the two involved parties (therapist vs. client) can be proved useful for identifying therapy-related behaviors.

Another direction for potential future improvements is related to the modeling approach followed for the utterance-level codes. The system presented here treats all the codes evenly and employs a single neural architecture giving one output label for every utterance. However, since human coders often stack multiple codes for a single utterance (e.g., asking for permission to give advice [ADP] through a

closed question [QUC]), a hierarchical algorithm which differentiates between codes with increasing granularity and allows for multiple codes per utterance may be useful. In such a scenario, a hybrid method which uses the modeling strength of neural networks and at the same time exploits knowledge-based information distilled from the coding manuals and clinical practice, can potentially improve the robustness and increase the interpretability of the results. This strategy would particularly benefit codes where our system performed relatively poorly (e.g., MI-adherent [MIA] and MI-non-adherent [MIN] behaviors; Table 8), due to limited training examples or due to insufficient information captured just from available linguistic cues. Keeping in mind that psychotherapy is a dyadic interaction, incorporating contextual information from the client's neighboring utterances could also lead to performance improvements, especially for codes such as reflections (RES and REC) that depend semantically on client's language (Table 2).

Limitations imposed by the available number of training samples is a crucial aspect regarding any machine learning based model. Even though herein we present and use one of the largest available corpora constructed for the purpose of automatic behavioral coding, the performance of all the models involved is still critically dependent on the sample size. This is why we decided to use a lot of third-party sources, both for training the behavior code predictors and for any audio or language modeling needed for the transcription pipeline. Applying external datasets, however, was not possible for all the tasks. In particular, for the session-level code prediction, we only had the internal 188 labeled samples available and we, hence, decided to apply a cross-validation scheme with a statistical model (support vector regression) that does not require as much data as a more convoluted deep learning model to converge. In any

case, all the results were reported on evaluation sets not seen during training, while the distributions of the predicted codes (Figs. 3 and 4) suggest that those results are indicative of the performance on a much bigger dataset of therapy sessions.

An aspect of importance for our system is the quality assurance of the final evaluation report provided to the counselor. Being able to determine computational errors at an early stage and giving relevant warning messages to the user is an essential prerequisite before mental health practitioners trust computer-based tools and introduce them into clinical settings. We have already implemented several quality safeguards, with results indicating that they are towards the right direction. We are planning on implementing more confidence metrics, which take into account ASR and behavior coding results, apart from VAD and diarization. Human annotators can still be used for the sessions or parts of sessions for which confidence is low. Such manually annotated sessions can be a valuable source of information to be used for further adapting our algorithms. That way, we can introduce an active learning scenario where the system incrementally becomes more accurate and reliable.

Likewise, it is important that we have evaluation metrics both for the individual modules and for the end-to-end system. Standard metrics, such as the word error rate (WER) and the diarization error rate (DER) used in ASR and in diarization, respectively, are useful during modeling in order to have benchmarks and quantifiable areas of improvement. However, they do not necessarily reflect the transcript quality from a user's perspective (Silovsky et al., 2012) and they are not always representative of the performance with respect to semantics and to clinical impact (Miner et al., 2020). Qualitative surveys where experts share their opinions on the accuracy of the system output could assist highlighting specific areas of clinical importance on which the modeling efforts should focus.

We should here underline that our goal is to build a system that will not replace the human input, but will instead assist medical experts increasing efficiency and accuracy. Technology-based tools have seen a rapid rise in healthcare with applications ranging from safety surveillance and epidemiological data collection (Cowie et al., 2017) to clinical decision-making and treatment recommendations (Sutton et al., 2020). However, all those tools, and especially the ones focusing on conversational interactions, are not expected to replace care providers, but rather augment their capabilities (Gangadharaiah et al., 2020). In the psychotherapy domain, an automatic evaluation platform, like the one we presented, would offer opportunities for ongoing self-assessment and self-improvement and would open new discussions on the development of specific skills between professionals or

between trainees and supervisors. Additionally, even with a widespread usage of automatic psychotherapy evaluation systems, the community will still need skilled and objective behavioral coders, both for the evaluation and for the training of the systems, since any machine learning algorithm is only as good as the training data we provide (Caliskan et al., 2017).

In any case, it is essential that the users be adequately trained to understand the meaning of an automatically generated feedback and what the several scores represent. It has been reported that experienced counselors are more likely to be skeptical about the validity of their ratings (Hirsch., 2018), as opposed to new and young therapists who may be attracted by the lure of machine learning, even without being fully aware of how their performance-based scores are estimated. Technology-based systems have the potential to transform mental healthcare. Being receptive to such a transformation should not mean uncritically accepting any machine-generated results. In fact, well-intentioned skepticism and criticism will accelerate the research in the field and will lead to an incremental improvement of the relevant technologies.

## Open practices statement

The original data collected for this study consist of real-world therapist–client sessions recorded at the University Counseling Center (UCC) of a large public western university and have to remain within the UCC servers at all times for privacy reasons; thus they cannot be made publicly available. The psychotherapy data used from previous studies (Tollison et al., 2008; Baer et al., 2009; Krupski et al., 2012; Neighbors et al., 2012; Lee et al., 2013; Lee et al., 2014) for adaptation are also protected and not publicly available. The speech corpora used to train the ASR system are either freely available or provided through the Language Data Consortium (LDC) to members and non-members for a fee ([www ldc upenn edu](http://www ldc upenn edu)). In particular, Librispeech (Panayotov et al., 2015) ([www openslr org/12](http://www openslr org/12)), TED-LIUM (Rousseau et al., 2014) ([lium univ lemans fr/ted-lium2](http://lium univ lemans fr/ted-lium2)), and AMI (Carletta et al., 2005) ([groups inf ed ac uk/ami/corpus](http://groups inf ed ac uk/ami/corpus)) are freely available to the community; Fisher English (Cieri et al., 2004) (Part 1: LDC2004S13 and Part 2: LDC2005S13), ICSI Meeting Speech (Janin et al., 2003) (LDC2004S02), WSJ (Paul & Baker, 1992) (Part 1: LDC93S6A and Part 2: LDC94S13A), and 1997 HUB4 (Graff et al., 1997) (LDC98S71) are provided through LDC. The Counseling and Psychotherapy Transcripts (without accompanying audio) that were used for some of the language-based modeling can be accessed on request at [alexanderstreet.com/products/counseling-and-psychotherapy-transcripts-series](http://alexanderstreet.com/products/counseling-and-psychotherapy-transcripts-series).



Our models are trained on real-world, sensitive, and protected data. Thus, our trained models cannot be made publicly available. Acoustic feature extraction and acoustic modeling was performed using the Kaldi toolkit which is available at [github.com/kaldi-asr/kaldi](https://github.com/kaldi-asr/kaldi). The BeamformIt tool used for acoustic beamforming is available at [github.com/xanguera/BeamformIt](https://github.com/xanguera/BeamformIt). Language models were built using the SRILM toolkit, available at [www.speech.sri.com/projects/srilm](http://www.speech.sri.com/projects/srilm). The neural network used for utterance-level code prediction was built on TensorFlow ([www.tensorflow.org](http://www.tensorflow.org)), while the tf-idf/SVR framework used for session-level code prediction made use of the scikit-learn Python library ([scikit-learn.org/stable](http://scikit-learn.org/stable)). The md-eval tool, developed by the National Institute of Standards and Technology (NIST) and used for diarization evaluation, is available at [github.com/usnistgov/SCTK/tree/master/src/md-eval](https://github.com/usnistgov/SCTK/tree/master/src/md-eval). The estimation of Krippendorff's alpha ( $\alpha$ ) for inter-rater reliability was based on the implementation available at [github.com/pln-fing-udelar/fast-krippendorff](https://github.com/pln-fing-udelar/fast-krippendorff).

None of the experiments was preregistered.

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.3758/s13428-021-01623-4>.

**Acknowledgements** Funding was provided by the National Institutes of Health/National Institute on Alcohol Abuse and Alcoholism (R01 AA018673).

**Author Contributions** Shrikanth Narayanan, Panayiotis Georgiou, David C. Atkins, and Zac E. Imel conceived the idea of the described system. James Gibson developed a first prototype of the processing pipeline. Nikolaos Flemotomos, Victor R. Martinez, Zhuohao Chen, Karan Singla, Victor Ardulov, and Raghuvveer Peri worked on training, adapting, and evaluating the various individual modules and the end-to-end system. Michael J. Tanana helped in testing the system in real-world settings, which led to several revisions and improvements. Jake Van Epps played a crucial role on data collection. Derek D. Caperton and Sarah P. Lord led the efforts in human behavioral coding. Tad Hirsch led the efforts of designing the graphical interface. Nikolaos Flemotomos reviewed the literature and drafted the original manuscript. All authors provided critical input, discussion, and revision of the manuscript and approved the final submitted version.

## Declarations

**Conflict of interests** Michael J. Tanana, Tad Hirsch, Zac E. Imel, David C. Atkins, and Shrikanth Narayanan are co-founders with equity stake in a technology company, Lyssn.io, focused on tools to support training, supervision, and quality assurance of psychotherapy and counseling. Sarah P. Lord also holds an equity stake in Lyssn.io. Shrikanth Narayanan is Chief Scientist and co-founder with equity stake in Behavioral Signal Technologies, a company focused on creating technologies for emotional and behavioral machine intelligence. James Gibson is a Senior Machine Learning Engineer in Behavioral Signal Technologies. The remaining authors report no conflicts of interest.

## References

- Anguera, X., Bozonnet, S., Evans, N., Fredouille, C., Friedland, G., & Vinyals, O. (2012). Speaker diarization: A review of recent research. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(2), 356–370.
- Anguera, X., Wooters, C., & Hernando, J. (2007). Acoustic beamforming for speaker diarization of meetings. *IEEE Transactions on Audio, Speech, and Language Processing*, 15(7), 2011–2022.
- Baer, J. S., Wells, E. A., Rosengren, D. B., Hartzler, B., Beadnell, B., & Dunn, C. (2009). Agency context and tailored training in technology transfer: A pilot evaluation of motivational interviewing training for community counselors. *Journal of Substance Abuse Treatment*, 37(2), 191–202.
- Bakeman, R., & Quera, V. (2012). Behavioral observation. In Cooper, H., Camic, P. M., Long, D. L., Panter, A. T., Rindskopf, D., & Sher, K. J. (Eds.) *APA handbook of research methods in psychology, foundations, planning, measures, and psychometrics*, (Vol. 1, pp. 207–225). Washington: American Psychological Association.
- Black, M. P., Katsamanis, A., Baucom, B. R., Lee, C.-C., Lammert, A. C., Christensen, A., et al. (2013). Toward automating a human behavioral coding system for married couples' interactions using speech acoustic features. *Speech Communication*, 55(1), 1–21.
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183–186.
- Can, D., Atkins, D. C., & Narayanan, S. S. (2015). A dialog act tagging approach to behavioral coding: A case study of addiction counseling conversations. In *Proceedings of annual conference of the international speech communication association* (pp. 339–343).
- Carletta, J., Ashby, S., Bourban, S., Flynn, M., Guillemot, M., Hain, T., et al. (2005). THE AMI meeting corpus: A pre-announcement. In *Proceedings of international workshop on machine learning for multimodal interaction* (pp. 28–39).
- Chen, Z., Flemotomos, N., Ardulov, V., Creed, T. A., Imel, Z. E., Atkins, D. C., & Narayanan, S. (2020). Feature fusion strategies for end-to-end evaluation of cognitive behavior therapy sessions. preprint at arXiv:2005.07809.
- Cieri, C., Miller, D., & Walker, K. (2004). The Fisher corpus: a resource for the next generations of speech-to-text. In *Proceedings of language resources and evaluation conference* (pp. 69–71).
- Cowie, M. R., Blomster, J. I., Curtis, L. H., Duclaux, S., Ford, I., Fritz, F., et al. (2017). Electronic health records to facilitate clinical research. *Clinical Research in Cardiology*, 106(1), 1–9.
- Curran, J., Parry, G. D., Hardy, G. E., Darling, J., Mason, A.-M., & Chambers, E. (2019). How does therapy harm? A model of adverse process using task analysis in the meta-synthesis of service users' experience. *Frontiers in Psychology*, 10, 347.
- Cutler, A., Dahan, D., & Van Donselaar, W. (1997). Prosody in the comprehension of spoken language: A literature review. *Language and Speech*, 40(2), 141–201.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of conference of the North American chapter of the association for computational linguistics: Human language technologies*. long and short papers, (Vol. 1, pp. 4171–4186).
- Flemotomos, N., Georgiou, P., & Narayanan, S. (2020). Linguistically aided speaker diarization using speaker role information. In *Proceedings of odyssey: The speaker and language recognition workshop* (pp. 117–124).
- Flemotomos, N., Martinez, V. R., Gibson, J., Atkins, D. C., Creed, T., & Narayanan, S. (2018). Language features for automated evaluation of cognitive behavior psychotherapy sessions. In

- Proceedings of annual conference of the international speech communication association* (pp. 1908–1912).
- Flemotomos, N., Papadopoulos, P., Gibson, J., & Narayanan, S. (2018). Combined speaker clustering and role recognition in conversational speech. In *Proceedings of annual conference of the international speech communication association* (pp. 1378–1382).
- Gangadharaiyah, R., Shivade, C., Bhatia, P., Zhang, Y., & Kass-Hout, T. (2020). Why conversational AI won't replace healthcare providers. In *Conversational agents for health and wellbeing, chi workshop*.
- Gaume, J., Gmel, G., Faouzi, M., & Daepfen, J.-B. (2009). Counselor skill influences outcomes of brief motivational interventions. *Journal of Substance Abuse Treatment*, 37(2), 151–159.
- Georgiou, P. G., Black, M. P., Lammert, A. C., Baucom, B. R., & Narayanan, S. S. (2011). That's aggravating, very aggravating: Is it possible to classify behaviors in couple interactions using automatically derived lexical features? In *International conference on affective computing and intelligent interaction* (pp. 87–96).
- Gibson, J., Atkins, D., Creed, T., Imel, Z., Georgiou, P., & Narayanan, S. (2019). Multi-label multi-task deep learning for behavioral coding. *IEEE Transactions on Affective Computing*.
- Goldberg, S. B., Flemotomos, N., Martinez, V. R., Tanana, M. J., Kuo, P. B., Pace, B. T., et al. (2020). Machine learning and natural language processing in psychotherapy research: Alliance as example use case. *Journal of Counseling Psychology*, 67(4), 438–448.
- Graff, D., Wu, Z., MacIntyre, R., & Liberman, M. (1997). The 1996 broadcast news speech and language-model corpus. In *Proceedings of DARPA workshop on spoken language technology* (pp. 11–14).
- Hallgren, K. A. (2012). Computing inter-rater reliability for observational data: an overview and tutorial. *Tutorials in quantitative methods for psychology*, 8(1), 23–34.
- Hattie, J., & Timperley, H. (2007). The power of feedback. *Review of educational research*, 77(1), 81–112.
- Heldner, M., & Edlund, J. (2010). Pauses, gaps and overlaps in conversations. *Journal of Phonetics*, 38(4), 555–568.
- Hill, C. E. (2009). *Helping skills: Facilitating, exploration, insight, and action*. Washington, DC: American Psychological Association.
- Hirsch, T., Soma, C., Merced, K., Kuo, P., Dembe, A., Caperton, D. D., et al. (2018). It's hard to argue with a computer<sup>2</sup>: Investigating psychotherapists' attitudes towards automated evaluation. In *Proceedings of designing interactive systems conference* (pp. 559–571).
- Houck, J. M., Moyers, T. B., Miller, W. R., Glynn, L. H., & Hallgren, K. A. (2010). Motivational interviewing skill code (misc) version 2.5. (Available from <http://casaa.unm.edu/download/misc25.pdf>).
- Imel, Z. E., Pace, B. T., Soma, C. S., Tanana, M., Hirsch, T., Gibson, J., et al. (2019). Design feasibility of an automated, machine-learning based feedback system for motivational interviewing. *Psychotherapy*, 56(2), 318.
- Imel, Z. E., Steyvers, M., & Atkins, D. C. (2015). Computational psychotherapy research: Scaling up the evaluation of patient-provider interactions. *Psychotherapy*, 52(1), 19.
- Ioffe, S. (2006). Probabilistic linear discriminant analysis. In *Proceedings of European conference on computer vision* (pp. 531–542).
- Janin, A., Baron, D., Edwards, J., Ellis, D., Gelbart, D., Morgan, N., et al. (2003). The ICSI meeting corpus. In *Proceedings of international conference on acoustics, speech, and signal processing* pp. 1–1.
- Klatte, R., Strauss, B., Flückiger, C., & Rosendahl, J. (2018). Adverse effects of psychotherapy: protocol for a systematic review and meta-analysis. *Systematic reviews*, 7(1), 135.
- Ko, T., Peddinti, V., Povey, D., & Khudanpur, S. (2015). Audio augmentation for speech recognition. In *Proceedings of annual conference of the international speech communication association* pp. 3586–3589.
- Kodish-Wachs, J., Agassi, E., Kenny, III, P., & Overhage, J. M. (2018). A systematic comparison of contemporary automatic speech recognition engines for conversational clinical speech. *Proc AMIA annual symposium, 2018*, 683.
- Krippendorff, K. (2018). *Content analysis: An introduction to its methodology*. Los Angeles, CA: Sage publications.
- Krupski, A., Joesch, J. M., Dunn, C., Donovan, D., Bumgardner, K., Lord, S. P., et al. (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.
- Kulik, J. A., & Kulik, C.-L. C. (1988). Timing of feedback and verbal learning. *Review of educational research*, 58(1), 79–97.
- Lambert, M. J., & Bergin, A. E. (2002). The effectiveness of psychotherapy. In Hersen, M., & Sledge, W. (Eds.) *Encyclopedia of psychotherapy*, (Vol. 1, pp. 709–714). USA: Elsevier Science.
- Lambert, M. J., & Ogles, B. M. (1997). The effectiveness of psychotherapy supervision. In Watkins, C. E. (Ed.) *Handbook of psychotherapy supervision*, (pp. 421–446). USA.
- Lambert, M. J., Whipple, J. L., & Kleinstäuber, M. (2018). Collecting and delivering progress feedback: A meta-analysis of routine outcome monitoring. *Psychotherapy*, 55(4), 520.
- Lee, C. M., Kilmer, J. R., Neighbors, C., Atkins, D. C., Zheng, C., Walker, D. D., & Larimer, M. E. (2013). Indicated prevention for college student marijuana use: A randomized controlled trial. *Journal of consulting and clinical psychology*, 81(4), 702.
- Lee, C. M., Neighbors, C., Lewis, M. A., Kaysen, D., Mittmann, A., Geisner, I. M., et al. (2014). Randomized controlled trial of a spring break intervention to reduce high-risk drinking. *Journal of consulting and clinical psychology*, 82(2), 189.
- Lee, F.-T., Hull, D., Levine, J., Ray, B., & McKeown, K. (2019). Identifying therapist conversational actions across diverse psychotherapeutic approaches. In *Proceedings of workshop on computational linguistics and clinical psychology* (pp. 12–23).
- Levitt, H. M. (2001). Sounds of silence in psychotherapy: The categorization of clients' pauses. *Psychotherapy Research*, 11(3), 295–309.
- Lord, S. P., Sheng, E., Imel, Z. E., Baer, J., & Atkins, D. C. (2015). More than reflections: empathy in motivational interviewing includes language style synchrony between therapist and client. *Behavior therapy*, 46(3), 296–303.
- Ma, X., & Hovy, E. (2016). End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF. *Proceedings of annual meeting of the association for computational linguistics*, 1 (1064–1074).
- Madson, M. B., & Campbell, T. C. (2006). Measures of fidelity in motivational enhancement: a systematic review. *Journal of substance abuse treatment*, 31(1), 67–73.
- Magill, M., Gaume, J., Apodaca, T. R., Walthers, J., Mastroleo, N. R., Borsari, B., & Longabaugh, R. (2014). The technical hypothesis of motivational interviewing: A meta-analysis of MI's key causal model. *Journal of consulting and clinical psychology*, 82(6), 973.
- Malik, U., Barange, M., Saunier, J., & Pauchet, A. (2018). Performance comparison of machine learning models trained on manual vs ASR transcriptions for dialogue act annotation. In *2018 IEEE 30th international conference on tools with artificial intelligence (ICTAI)* (pp. 1013–1017).
- Martinez, V. R., Flemotomos, N., Ardulov, V., Somandepalli, K., Goldberg, S. B., Imel, Z. E., et al. (2019). Identifying therapist and client personae for therapeutic alliance estimation. *Proceedings of*

- Annual Conference of the International Speech Communication Association* (pp. 1901–1905).
- Miller, W. R., & Rollnick, S. (2012). *Motivational interviewing: Helping people change*. Guilford press.
- Miller, W. R., Sorensen, J. L., Selzer, J. A., & Brigham, G. S. (2006). Disseminating evidence-based practices in substance abuse treatment: A review with suggestions. *Journal of substance abuse treatment, 31*(1), 25–39.
- Miner, A. S., Haque, A., Fries, J. A., Fleming, S. L., Wilfley, D. E., Wilson, G. T., et al. (2020). Assessing the accuracy of automatic speech recognition for psychotherapy. *NPJ Digital Medicine, 3*(82), 82.
- Moyers, T. B., Martin, T., Manuel, J. K., Hendrickson, S. M., & Miller, W. R. (2005). Assessing competence in the use of motivational interviewing. *Journal of substance abuse treatment, 28*(1), 19–26.
- Moyers, T. B., Rowell, L. N., Manuel, J. K., Ernst, D., & Houck, J. M. (2016). The motivational interviewing treatment integrity code (MITI 4): rationale, preliminary reliability and validity. *Journal of substance abuse treatment, 65*, 36–42.
- Nasir, M., Chakravarthula, S. N., Baucom, B. R., Atkins, D. C., Georgiou, P., & Narayanan, S. (2019). Modeling interpersonal linguistic coordination in conversations using word mover's distance. Proceedings of Annual Conference of the International Speech Communication Association (pp. 1423–1427).
- Neighbors, C., Lee, C. M., Atkins, D. C., Lewis, M. A., Kaysen, D., Mittmann, A., et al. (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*(5), 850.
- Panayotov, V., Chen, G., Povey, D., & Khudanpur, S. (2015). Librispeech: an ASR corpus based on public domain audio books. In *Proceedings of international conference on acoustics, speech and signal processing* (pp. 5206–5210).
- Paul, D. B., & Baker, J. M. (1992). The design for the Wall Street Journal-based CSR corpus. In *Proceedings of workshop on speech and natural language* (pp. 357–362).
- Peddinti, V., Chen, G., Manohar, V., Ko, T., Povey, D., & Khudanpur, S. (2015). JHU ASPIRE system: Robust LVCSR with TDNNs, i-vector adaptation and RNN-LMs. In *Proceedings of workshop on automatic speech recognition and understanding* (pp. 539–546).
- Peddinti, V., Povey, D., & Khudanpur, S. (2015). A time delay neural network architecture for efficient modeling of long temporal contexts. In *Proceedings of annual conference of the international speech communication association* (pp. 3214–3218).
- Perry, J. C., Banon, E., & Ianni, F. (1999). Effectiveness of psychotherapy for personality disorders. *American Journal of Psychiatry, 156*(9), 1312–1321.
- Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., et al. (2011). The Kaldi Speech Recognition Toolkit. In *Proceedings of workshop on automatic speech recognition and understanding*.
- Prince, S. J., & Elder, J. H. (2007). Probabilistic linear discriminant analysis for inferences about identity. In *Proceedings of international conference on computer vision* (pp. 1–8).
- Proctor, E., Silmere, H., Raghavan, R., Hovmand, P., Aarons, G., Bunker, A., & Hensley, M. (2011). Outcomes for implementation research: conceptual distinctions, measurement challenges, and research agenda. *Administration and Policy in Mental Health and Mental Health Services Research, 38*(2), 65–76.
- Quiroz, J. C., Laranjo, L., Kocaballi, A. B., Berkovsky, S., Rezazadegan, D., & Coiera, E. (2019). Challenges of developing a digital scribe to reduce clinical documentation burden. *npj Digital Medicine, 2*(1), 1–6.
- Rojas-Barahona, L., Tseng, B.-H., Dai, Y., Mansfield, C., Ramadan, O., Ultes, S., et al. (2018). Deep learning for language understanding of mental health concepts derived from cognitive behavioural therapy. In *Proceedings of international workshop on health text mining and information analysis* (pp. 44–54).
- Rousseau, A., Delüglise, P., & Esteve, Y. (2014). Enhancing the TED-LIUM corpus with selected data for language modeling and more TED talks. In *Proceedings of language resources and evaluation conference* (pp. 3935–3939).
- Salton, G., & McGill, M. J. (1986). *Introduction to modern information retrieval*. New York: McGraw-Hill, Inc.
- Saxon, D., Barkham, M., Foster, A., & Parry, G. (2017). The contribution of therapist effects to patient dropout and deterioration in the psychological therapies. *Clinical psychology & psychotherapy, 24*(3), 575–588.
- Schmidt, L. K., Andersen, K., Nielsen, A. S., & Moyers, T. B. (2019). Lessons learned from measuring fidelity with the motivational interviewing treatment integrity code (MITI 4). *Journal of Substance Abuse Treatment, 97*, 59–67.
- Schwalbe, C. S., Oh, H. Y., & Zweben, A. (2014). Sustaining motivational interviewing: A meta-analysis of training studies. *Addiction, 109*(8), 1287–1294.
- Sell, G., Snyder, D., McCree, A., Garcia-Romero, D., Villalba, J., Maciejewski, M., et al. (2018). Diarization is hard: Some experiences and lessons learned for the JHU team in the inaugural DIHARD challenge. In *Proceedings of annual conference of the international speech communication association* (pp. 2808–2812).
- Shiner, B., D'Avolio, L. W., Nguyen, T. M., Zayed, M. H., Watts, B. V., & Fiore, L. (2012). Automated classification of psychotherapy note text: implications for quality assessment in PTSD care. *Journal of evaluation in clinical practice, 18*(3), 698–701.
- Silovsky, J., Zdansky, J., Nouza, J., Cerva, P., & Prazak, J. (2012). Incorporation of the ASR output in speaker segmentation and clustering within the task of speaker diarization of broadcast streams. In *Proceedings of international workshop on multimedia signal processing* (pp. 118–123).
- Singla, K., Chen, Z., Atkins, D. C., & Narayanan, S. (2020). Towards end-2-end learning for predicting behavior codes from spoken utterances in psychotherapy conversations. In *Proceedings of annual meeting of the association for computational linguistics* (pp. 3797–3803).
- Singla, K., Chen, Z., Flemotomos, N., Gibson, J., Can, D., Atkins, D. C., & Narayanan, S. (2018). Using prosodic and lexical information for learning utterance-level behaviors in psychotherapy. In *Proceedings of annual conference of the international speech communication association* (pp. 3413–3417).
- Snyder, D., Garcia-Romero, D., Sell, G., Povey, D., & Khudanpur, S. (2018). X-vectors: Robust DNN embeddings for speaker recognition. In *Proceedings of international conference on acoustics, speech and signal processing* (pp. 5329–5333).
- Stolcke, A. (2002). SRILM—an extensible language modeling toolkit. In *Proceedings of international conference on spoken language processing* (pp. 901–904).
- Sutton, R. T., Pincok, D., Baumgart, D. C., Sadowski, D. C., Fedorak, R. N., & Kroeker, K. I. (2020). An overview of clinical decision support systems: benefits, risks, and strategies for success. *NPJ Digital Medicine, 3*(1), 1–10.
- Thomas, S., Saon, G., Van Segbroeck, M., & Narayanan, S. S. (2015). Improvements to the IBM speech activity detection system for the DARPA RATS program. In *Proceedings of international conference on acoustics, speech and signal processing* (pp. 4500–4504).
- Tollison, S. J., Lee, C. M., Neighbors, C., Neil, T. A., Olson, N. D., & Larimer, M. E. (2008). Questions and reflections: the use of motivational interviewing microskills in a peer-led brief alcohol intervention for college students. *Behavior Therapy, 39*(2), 183–194.

- Weisz, J. R., Weiss, B., Han, S. S., Granger, D. A., & Morton, T. (1995). Effects of psychotherapy with children and adolescents revisited: a meta-analysis of treatment outcome studies. *Psychological Bulletin*, *117*(3), 450.
- Xiao, B., Bone, D., Segbroeck, M. V., Imel, Z. E., Atkins, D. C., Georgiou, P. G., & Narayanan, S. S. (2014). Modeling therapist empathy through prosody in drug addiction counseling. In *Proceedings of annual conference of the international speech communication association* (pp. 213–217).
- Xiao, B., Can, D., Georgiou, P. G., Atkins, D., & Narayanan, S. S. (2012). Analyzing the language of therapist empathy in motivational interview based psychotherapy. In *Proceedings of Asia Pacific signal and information processing association annual summit and conference* (pp. 1–4).
- Xiao, B., Can, D., Gibson, J., Imel, Z. E., Atkins, D. C., Georgiou, P. G., & Narayanan, S. S. (2016). Behavioral coding of therapist language in addiction counseling using recurrent neural networks. In *Proceedings of annual conference of the international speech communication association* (pp. 908–912).
- Xiao, B., Huang, C., Imel, Z. E., Atkins, D. C., Georgiou, P., & Narayanan, S. S. (2016). A technology prototype system for rating therapist empathy from audio recordings in addiction counseling. *PeerJ Computer Science*, *2*, e59.
- Xiao, B., Imel, Z. E., Atkins, D. C., Georgiou, P. G., & Narayanan, S. S. (2015). Analyzing speech rate entrainment and its relation to therapist empathy in drug addiction counseling. In *Proceedings of annual conference of the international speech communication association* (pp. 2489–2493).
- Xiao, B., Imel, Z. E., Georgiou, P. G., Atkins, D. C., & Narayanan, S. S. (2015). “Rate my therapist”: Automated detection of empathy in drug and alcohol counseling via speech and language processing. *PLoS one*, *10*(12).
- Xiong, W., Droppo, J., Huang, X., Seide, F., Seltzer, M. L., Stolcke, A., et al. (2017). Toward human parity in conversational speech recognition. *IEEE/ACM Transactions on Audio Speech, and Language Processing*, *25*(12), 2410–2423.

**Publisher’s note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## Affiliations

Nikolaos Flemotomos<sup>1</sup> · Victor R. Martinez<sup>2</sup> · Zhuohao Chen<sup>1</sup> · Karan Singla<sup>2</sup> · Victor Ardulov<sup>2</sup> · Raghuvver Peri<sup>1</sup> · Derek D. Caperton<sup>3</sup> · James Gibson<sup>4</sup> · Michael J. Tanana<sup>5</sup> · Panayiotis Georgiou<sup>1</sup> · Jake Van Epps<sup>6</sup> · Sarah P. Lord<sup>7</sup> · Tad Hirsch<sup>8</sup> · Zac E. Imel<sup>3</sup> · David C. Atkins<sup>7</sup> · Shrikanth Narayanan<sup>1,2,4</sup>

<sup>1</sup> Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, California, USA

<sup>2</sup> Department of Computer Science, University of Southern California, Los Angeles, CA, 90089, USA

<sup>3</sup> Department of Educational Psychology, University of Utah, Salt Lake City, Utah, USA

<sup>4</sup> Behavioral Signal Technologies Inc., Los Angeles, CA, USA

<sup>5</sup> College of Social Work, University of Utah, Salt Lake City, Utah, USA

<sup>6</sup> University Counseling Center, University of Utah, Salt Lake City, Utah, USA

<sup>7</sup> Department of Psychiatry and Behavioral Sciences, University of Washington, Seattle, Washington, USA

<sup>8</sup> Department of Art + Design, Northeastern University, Boston, Massachusetts, USA