


METHODOLOGY

Open Access



# Heterogeneous separation consistency training for adaptation of unsupervised speech separation

Jiangyu Han<sup>1,2</sup> and Yanhua Long<sup>1,2\*</sup> 

## Abstract

Recently, supervised speech separation has made great progress. However, limited by the nature of supervised training, most existing separation methods require ground-truth sources and are trained on synthetic datasets. This ground-truth reliance is problematic, because the ground-truth signals are usually unavailable in real conditions. Moreover, in many industry scenarios, the real acoustic characteristics deviate far from the ones in simulated datasets. Therefore, the performance usually degrades significantly when applying the supervised speech separation models to real applications. To address these problems, in this study, we propose a novel separation consistency training, termed SCT, to exploit the real-world unlabeled mixtures for improving cross-domain unsupervised speech separation in an iterative manner, by leveraging upon the complementary information obtained from heterogeneous (structurally distinct but behaviorally complementary) models. SCT follows a framework using two heterogeneous neural networks (HNNs) to produce high confidence pseudo labels of unlabeled real speech mixtures. These labels are then updated and used to refine the HNNs to produce more reliable consistent separation results for real mixture pseudo-labeling. To maximally utilize the large complementary information between different separation networks, a cross-knowledge adaptation is further proposed. Together with simulated dataset, those real mixtures with high confidence pseudo labels are then used to update the HNN separation models iteratively. In addition, we find that combining the heterogeneous separation outputs by a simple linear fusion can further slightly improve the final system performance. In this paper, we use cross-dataset to simulate the cross-domain situation in real-life. The term of “source domain” and “target domain” refer to the simulation set for model pre-training and the real unlabeled mixture for model adaptation. The proposed SCT is evaluated on both public reverberant English and anechoic Mandarin cross-domain separation tasks. Results show that, without any available ground-truth of target domain mixtures, the SCT can still significantly outperform our two strong baselines with up to 1.61 dB and 3.44 dB scale-invariant signal-to-noise ratio (SI-SNR) improvements, on the English and Mandarin cross-domain conditions, respectively.

**Keywords** Unsupervised speech separation, Heterogeneous, Separation consistency, Cross-knowledge adaptation

## 1 Introduction

Multi-speaker interaction scenarios are very common in real-world speech processing applications. Speech separation, separating each source signal from mixed speech, is one of the most important technology for these applications, including speaker diarization, speaker verification, multi-talker speech recognition, etc.

Because of the importance of speech separation, numerous studies have focused on this topic, including

\*Correspondence:

Yanhua Long  
yanhua@shnu.edu.cn

<sup>1</sup> Key Innovation Group of Digital Humanities Resource and Research, Shanghai Normal University, Shanghai, China

<sup>2</sup> Shanghai Engineering Research Center of Intelligent Education and Bigdata, Shanghai Normal University, Shanghai, China

the traditional time-frequency (T-F) domain separation methods [1–12], and the recent popular time-domain approaches [13–20]. All these contributions have led to significant progress on the single-channel speech separation. Most of them follow a mask learning pattern that aims to learn a weighting matrix (mask) to capture relationship between the isolated clean sources. The mask is then used to separate each source signal with an element-by-element multiplication. In addition, some researchers also concentrate on learning clean sources directly from the mixed speech, which is known as mapping-based separation [21–23].

Reviewing recent speech separation techniques, most of them are supervised ones with their own advantages. Such as, the T-F domain methods take spectrogram as input features that are good at capturing the phonetic structure of speech [24]; the time-domain methods pay more attention to the fine structure of speech but are vulnerable to environmental or background variations; the masking-based methods are effective in utilizing the clean speech of training corpus but sensitive to speech with signal-to-noise ratio (SNR) variations; the mapping-based methods show more robustness for tasks with a wide range of SNR [25], etc. To fully exploit advantages over different approaches, some studies focus on integrating different approaches into an ensemble training framework. For example, authors in [24] constructed a time-and-frequency feature map by concatenating both time and time-frequency domain acoustic features to improve separation performance. For improving the singing voice extraction, in [26], several attention-based fusion strategies were proposed to utilize the complementarities between masking and mapping spectrograms using a minimum difference masks (MDMs) [27] criterion.

Although the supervised speech separation methods or their combinations have performed well on data with the same or similar acoustic properties as the simulation training sets, the performance on cross-domain real-world mixtures is still quite poor. The main problem of supervised training is the strong reliance on individual ground-truth source signals. It heavily precludes technique scaling to widely available real-world mixtures, and limits progress on wide-domain coverage speech separation tasks. In real scenarios, the isolated sources are difficult to collect, because they are usually contaminated by cross-talk and unknown acoustic channel impulse responses. Therefore, it's very difficult to provide golden-standard handcrafted labels for a large number of real-world mixtures to train a supervised separation system from scratch. Moreover, adapting a well pre-trained supervised system to target real acoustics is also challenging, because the distribution

of sound types and reverberation may be unknown and hard to estimate.

One approach to improve real-world unsupervised speech separation is to directly use the real acoustics in system training. To this end, some latest works start to separate speech from unsupervised or semi-supervised perspectives. In [28–30], a mixture invariant training (MixIT) that requires only single-channel real acoustic mixtures was proposed. MixIT uses mixtures of mixtures (MoMs) as input, and sums over estimated sources to match the target mixtures instead of the single-source references. As the model is trained to separate the MOMs into a variable number of latent sources, the separated sources can be remixed to approximate the original mixtures. Motivated by MixIT, authors in [31] proposed a teacher-student MixIT (TS-MixIT) to alleviate the over-separation problem in the original MixIT. It takes the unsupervised model trained by MixIT as a teacher model, then the estimated sources are filtered and selected as pseudo-targets to further train a student model using standard permutation invariant training (PIT) [3]. Besides, there are other unsupervised separation attempts as well, such as the co-separation [32], adversarial unmix-and-remix [33], and Mixup-Breakdown [34]. All these recent efforts indicate how to well exploit the real-world unlabeled mixtures to boost the current separation systems becomes very fundamental, important, and challenging.

In this study, we also focus on improving the unsupervised speech separation, a novel speech separation adaptation framework, termed separation consistency training (SCT), is proposed. Different from previous works, SCT aims to leverage the separation consistency between heterogeneous separation networks to produce high confidence pseudo-labels of unlabeled acoustic mixtures. These labels and networks are updated iteratively using a cross-knowledge adaptation approach to achieve more accurate pseudo-labels and better target speech separation models. In SCT, two separation networks with a heterogeneous structure are used, one is the current popular masking-based time-domain speech separation model, Conv-TasNet [13], and the other is our recent proposed mapping-based time-frequency domain separation model, DPCCN [35]. These two networks are then used to generate consistent separation results for target domain unlabeled mixture labeling. Considering the high freedom and strong modeling ability in deep neural networks, each heterogeneous model with a consistent learning objective tends to achieve similar performance in its own distinct way. Correspondingly, the difference between the separation models would result in a discriminative but complementary behavior to separate speech. Although this behavior is somehow vague for

humans due to the “black-box” property of neural networks. Therefore, the advantages behind using heterogeneous networks instead of homogeneous ones are that, besides the mixture labeling, the complementary information between these heterogeneous models is expected to attain large diversity for label creation. By doing so, it provides an increased chance to produce and select more informative target mixtures as iterative training samples that a single source separation model could not produce by itself. In addition, a simple linear fusion strategy is proposed to combine the heterogeneous separation outputs to further improve the final separation performance.

Our experiments are performed on three open-source datasets, the anechoic English Libri2Mix [36] is taken as the source domain data, the reverberant English WHAMR! [37] and anechoic Mandarin Aishell2Mix [35] are our two target domain datasets. Extensive results show that, the proposed SCT is very effective to improve the unsupervised cross-domain speech separation performance. It can significantly outperform two strong baselines with up to 1.61 dB and 3.44 dB scale-invariant signal-to-noise ratio (SI-SNR) [38] improvements on the English and Mandarin cross-domain tasks, respectively. Besides, we find that, our separation consistency selection can achieve competitive performance with the data selection using ground-truth sources as references during the target heterogeneous model adaptation.

## 2 Previous work

### 2.1 Conv-TasNet

Conv-TasNet is a time-domain, masking-based speech separation technique that proposed in [13]. Compared with most time-frequency domain algorithms, Conv-TasNet shows superior separation performance on the standard public WSJ0-2mix [1] dataset, and has become the mainstream speech separation approach. This network has attracted widespread attention and been further improved in many recent works [39–42].

Conv-TasNet consists of three parts: an encoder (1d convolution layer), a mask estimator (several convolution blocks), and a decoder (1d deconvolution layer). The waveform mixture is first encoded by the encoder and then is fed into the temporal convolutional network (TCN) [43–45] based mask estimator to estimate a multiplicative masking function for each source. Finally, the source waveforms are reconstructed by transforming the masked encoder representations using the decoder. More details can be found in [13].

### 2.2 DPCCN

DPCCN is our recent work in [35], it is a time-frequency domain, mapping-based speech separation technique. Results in [35] show that DPCCN can achieve much

better performance and robustness over other state-of-the-art separation methods in environmental complicated conditions.

DPCCN follows a U-Net [46] style to encode the mixture spectrum into a high-level representation, then decodes it into the clean speech. In DPCCN, DenseNet [47] is used to alleviate the vanishing-gradient problem and encourage the feature reuse; TCN is clamped between the codec to leverage long-range time information; A pyramid pooling layer [48] is introduced after decoder to improve its global modeling ability. The detailed information can be found in [35].

## 3 Heterogeneous separation consistency training

The proposed separation consistency training is performed on two different separation networks with heterogeneous structure. In this section, we first present the principle of our SCT, then introduce three SCT variants and their differences, including basic SCT and the cross-knowledge adaptation. Next, two main algorithms, consistent pseudo-labeling and selection (CPS), and heterogeneous knowledge fusion (HKF) in the proposed SCT are described in detail. For simplicity, here we only consider the speech separation scenario with two-speaker overlapped speech.

### 3.1 Separation consistency training

Our separation consistency training is specially designed to improve the unsupervised speech separation where the target mixtures deviate far from the training simulation dataset. It follows a heterogeneous separation framework, to create and select informative data pairs with high confidence pseudo ground-truth, for iteratively improving cross-domain speech separation by adapting the source separation models to the target acoustic environments. Because the whole framework heavily relies on the consistent separation results of the unlabeled mixtures and a separation consistency measure for pseudo-labeling, we call the whole training process as separation consistency training (SCT).

#### 3.1.1 Basic SCT

Given a large amount of or even limited unlabeled target mixtures, the basic SCT procedure can be divided into three main steps:

- (a) *Mixture separation.* Separate each unlabeled mixture using two heterogeneous separation models that have been well-trained on the source simulated training set;
- (b) *Consistent pseudo-labeling and selection (CPS).* Based on separation results in step (a), calculate a separation consistency measure (SCM, Eq. (1)) and

a mixture separation consistency measure (mSCM, Eq. (3)) to evaluate the confidence of separation outputs. Then, select those unlabeled mixtures with high consistent confidence and their corresponding separation results as pseudo ground-truth to form a “Pseudo Labeled Set”;

- (c) *Iterative model adaptation.* Combine the “Pseudo Labeled Set” and the original source domain “Simulation Training Set” together to refine the source models to learn the target domain acoustics. And then, repeat the above process in an iterative manner.


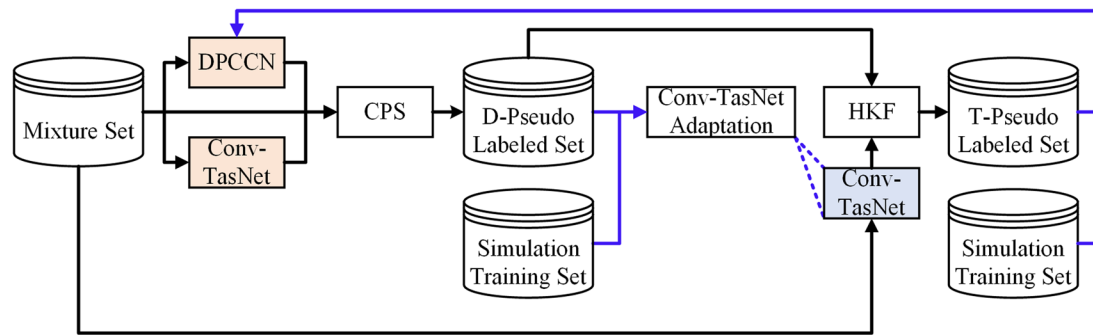
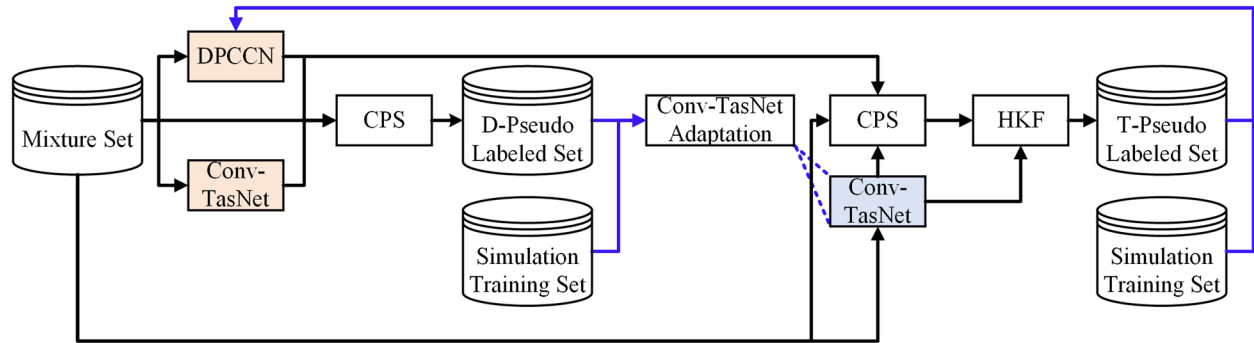
The two separation models in step (a) usually have comparable performance but with heterogeneous neural network structures. The bigger difference between the models, the better complementary information will be achieved. In this study, we choose DPCCN and Conv-TasNet that presented in Section 2 as the heterogeneous candidates. The former is taken as the primary model, while the latter is regarded as a reviewer model. Conv-TasNet [13] is the current popular masking-based time-domain separation model, while DPCCN [35] is our recent proposed mapping-based time-frequency domain model with good robustness to complicate acoustic environments. The huge difference in modeling patterns between Conv-TasNet and DPCCN, such as masking and mapping, time domain and time-frequency domain, guarantees the large diversity of the separated results. This diversity provides an increased chance to improve source models iteratively, because it can produce more informative target mixtures as new iterative training samples that the primary source model could not produce itself. Actually, during CPS in step (b), each model in this SCT heterogeneous framework is a reviewer for the other, any input mixtures will be double inspected by the two reviewers from different perspectives, only those mixtures with consistent separation performance of both will be selected. In this way, the double inspecting mechanism under a heterogeneous framework ensures the high confidence of pseudo ground-truth for each selected mixture in the target domain.

The whole framework of above basic SCT is demonstrated in the first variant of our proposed SCT, subfigure (A) (SCT-1) of Fig. 1. In SCT-1, the detail of consistent pseudo-labeling and selection (CPS) is presented in the next section, Section 3.2, and illustrated in Fig. 2A. “D-Pseudo Labeled Set” (“D-” means DPCCN’s outputs) contains the data pairs of selected unlabeled mixtures and their pseudo ground-truth that derive from the individual separation outputs of the primary model DPCCN. Together with the original source domain “Simulation Training Set,” both the primary and reviewer models

are refined and adapted to the target domain in each single iteration. It is worth noting that the model adaptation with the combined training set is necessary for SCT algorithm. As our source models have been trained well on the simulation set, and the pseudo ground-truth of “D-Pseudo Labeled Set” is actually generated by DPCCN, which means if we only use the simulation set or “D-Pseudo Labeled Set” to adjust the primary source model, DPCCN, the training gradient will be very small or even 0. In this case, the error between model outputs and labels is difficult to back propagate and the adaptation process will fail. However, if we adapt model using both “Simulation Training Set” and “D-Pseudo Labeled Set,” although the error between model outputs and ground-truth is small, the model can still be adapted to the target domain. For example, a simple neural network can be depicted as  $y = w * x + b$ , where  $y, w, x, b$  are model output, parameter weight, model input, and parameter bias, respectively. The partial differential to the weight  $w$  is model input  $x$ . Back to our scenario, by combining “Simulation Training Set” and “D-Pseudo Labeled Set,” the target domain data can engage in the model adaptation with the loss of the source domain simulation set.

### 3.1.2 SCT with cross-knowledge adaptation

To fully exploit the complementary information between heterogeneous networks, a cross-knowledge adaptation is proposed to improve the basic SCT. The framework is illustrated in the 2nd variant of SCT (SCT-2) in Fig. 1B. Different from basic SCT, in SCT-2, the reviewer Conv-TasNet is first updated using the combined “D-Pseudo Labeled Set” and “Simulation Training Set,” i.e., the pseudo ground-truth of the primary model is used to guide the reviewer model’s adaptation. Next, we re-separate all the unlabeled mixtures using the updated reviewer to achieve more accurate separation outputs. Then, all the pseudo ground-truth in “D-Pseudo Labeled Set” are replaced by the corresponding new individual outputs that produced by the updated reviewer Conv-TasNet to construct a new pseudo labeled set “T-Pseudo Labeled Set” (“T-” means Conv-TasNet’s outputs). Finally, the “T-Pseudo Labeled Set” and “Simulation Training Set” are combined together to adjust the primary model DPCCN as in SCT-1. In this model adaptation, the pseudo ground-truth of the reviewer model is used to supervise the primary model training. Just like the teacher-student learning, in the whole SCT-2, the primary and reviewer model can benefit each other, the learned knowledge of them is cross-used as a guide to improve the target model adaptation. Therefore, we call this adaptation procedure as “cross-knowledge adaptation” for simplicity. In addition, as the “T-Pseudo Labeled

**A. Separation consistency training-1 (SCT-1)****B. Separation consistency training-2 (SCT-2)****C. Separation consistency training-3 (SCT-3)**

**Fig. 1** The flowcharts of three SCT variants for single iteration. **A** Framework of the 1st SCT variant (SCT-1). The selected DPCCN outputs with pseudo labels are used to update both Conv-TasNet and DPCCN. **B** Framework of the 2nd SCT variant (SCT-2) with the cross-knowledge adaptation. **C** Framework of the 3rd SCT variant (SCT-3). Two-stage CPS are used to refine the pseudo-labeling

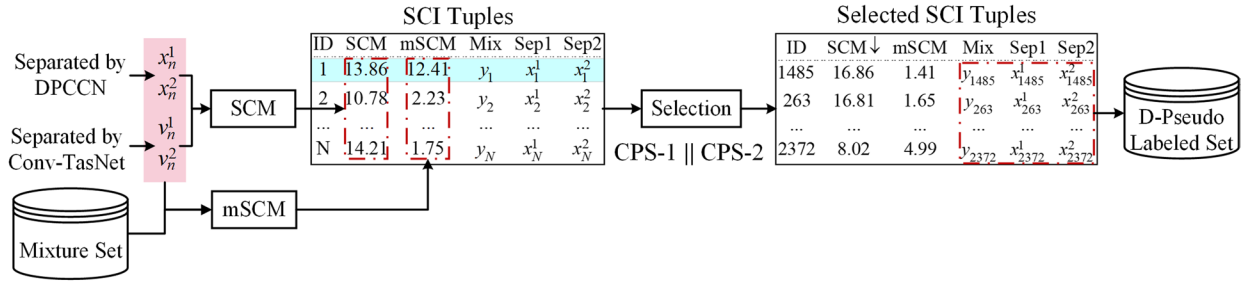
Set” is actually a combination of prior selected separation consistency statistics in “D-Pseudo Labeled Set” and the new pseudo ground-truth from updated Conv-TasNet, thus, in Fig. 1, we use the heterogeneous knowledge fusion (HKF) block to represent this knowledge combination. Details of HKF are demonstrated in subfigure (D) of Fig. 2 and Section 3.3.

Subfigure (C) (SCT-3) of Fig. 1 is a variant of SCT-2 with minor modification before HKF block. In SCT-3, the CPS is performed twice. The first CPS is performed as


the same in SCT-1 and SCT-2, while in the second CPS, the separation consistency statistics,  $SCM$  and  $mSCM$  are re-computed and updated using both mixture separation outputs of DPCCN and the updated Conv-TasNet. Other operations are all the same as in SCT-2. The motivation behind this two-stage CPS is that, the adapted Conv-TasNet can provide more accurate separation results of target domain mixtures, which makes the second stage CPS produce more reliable consistent separation results for unlabeled mixture pseudo-labeling in each SCT iteration.




### A. Consistent pseudo-labeling and selection (CPS)




### B. Separation consistency measure (SCM)



### C. Mixture separation consistency measure (mSCM)



### D. Heterogeneous knowledge fusion (HKF)



**Fig. 2** **A** The whole framework of consistent pseudo-labeling and selection (CPS). **B** The flowchart of separation consistency measure. **C** The flowchart of the mixture separation consistency measure. **D** The heterogeneous knowledge fusion

In summary, in this section, we present three variants of the proposed SCT, one is the basic structure, and the others are two enhanced SCT variants with cross-knowledge adaptation. Details of the CPS and HKF blocks used in SCT are described in the following sections.

#### 3.2 Consistent pseudo-labeling and selection

The consistent pseudo-labeling and selection (CPS) block in the proposed SCT aims to produce high quality individual pseudo ground-truth of each unlabeled mixture based on the outputs of two heterogeneous networks and the original mixture speech. The whole CPS procedure is illustrated Fig. 2A. It contains two main stages, the first one is the confidence measure calculation, followed by the pseudo ground-truth selection.

##### 3.2.1 Confidence measure calculation

Two measures are calculated in this stage, one is separation consistency measure (SCM, Eq. (1)), and the other is mixture separation consistency measure (mSCM, Eq. (3)). Both of them are used to evaluate the confidence

between heterogeneous separation outputs produced by DPCCN and Conv-TasNet.

As shown in the left part of Fig. 2A, given  $N$  unsupervised mixed speech with each contains  $M$  single sources, here we assume  $M = 2$ . For the  $n$ th mixture, the SCM is calculated by taking the individual separation output  $\mathbf{x}_n$  of the primary model DPCCN as pseudo reference as follows:

$$\text{SCM}(\mathbf{x}_n, \mathbf{v}_n) = \max_{\mathbf{P}} \frac{1}{M} \sum_{i=1}^M \text{SI-SNR}(\mathbf{x}_n^i, [\mathbf{P}\mathbf{v}_n]_i) \quad (1)$$

where  $\mathbf{x}_n = [x_n^1, x_n^2, \dots, x_n^M]^T$ ,  $\mathbf{v}_n = [v_n^1, v_n^2, \dots, v_n^M]^T$  are the  $M$  individual separation speech signals that separated by DPCCN and Conv-TasNet for the  $n$ -th input mixture, respectively.  $x_n^i$  and  $v_n^i$  are the  $i$ -th individual signal.  $\mathbf{P}$  is an  $M \times M$  permutation matrix,  $[\cdot]_i$  denotes selecting  $i$ -th element from the matrix, and  $\mathbf{T}$  is the operation of transpose. The SI-SNR in Eq. (1) is the standard scale-invariant signal-to-noise ratio (SI-SNR) [38] that used to measure the performance of state-of-the-art speech separation systems. It is defined as:

$$SI - SNR(s, \hat{s}) = 10 \log_{10} \left( \frac{\| \frac{\langle \hat{s}, s \rangle}{\langle s, s \rangle} s \|^2}{\| \frac{\langle \hat{s}, s \rangle}{\langle s, s \rangle} s - \hat{s} \|^2} \right) \quad (2)$$

where  $s$  and  $\hat{s}$  are the reference and estimated speech, respectively.  $\| \cdot \|^2$  denotes the signal power,  $\langle \cdot \rangle$  is the inner-product operation.

Figure 2B shows a two-speaker SCM process for the  $n$ -th mixture. The DPCCN outputs,  $x_n^1, x_n^2$  are taken as references to calculate the pairwise SI-SNR with the Conv-TasNet outputs,  $v_n^1$  and  $v_n^2$ . In this case, there are two permutation combinations, namely  $[x_n^1, v_n^1; x_n^2, v_n^2]$  and  $[x_n^1, v_n^2; x_n^2, v_n^1]$ , then SCM compares the averaging pairwise SI-SNR for each assignment and takes the highest value to represent the separation consistency between two heterogeneous networks outputs. The higher SCM, the better consistency of unlabeled separation outputs we can trust. However, when the input mixtures are hard to separate for both heterogeneous networks,  $\mathbf{x}_n$  and  $\mathbf{v}_n$  can be very close to the original mixture speech, and they could also result in a very high SCM. In this case, the pseudo reference  $\mathbf{x}_n$  may be far from the ground-truth and may not be qualified for the source model adaptation. To alleviate this situation, the following mSCM is introduced from another perspective to evaluate the quality of target domain mixture separation results and enhance the confidence of selected results.

The mixture separation consistency measure (mSCM), aims to measure the consistency between the outputs of heterogeneous networks and the original input mixture  $y_n$ . It is defined as:

$$mSCM(y_n, \mathbf{x}_n, \mathbf{v}_n) = \frac{1}{2M} \sum_{i=1}^M \sum_{\phi} SI - SNR(y_n, \phi_n^i) \quad (3)$$

where  $\phi_n^i \in \{x_n^i, v_n^i\}$  is the  $i$ -th individual output of DPCCN or Conv-TasNet of  $n$ -th input mixture as shown in Eq. (1). Figure 2C gives a detailed operation of mSCM under a two-speaker case, and as shown in Eq. (3), we see the average SI-SNR between the input mixture and all separated outputs are calculated. Different from SCM, the mSCM evaluate the confidence of separation results in an opposite way and the lower is desired. We believe that, in most conditions, the waveform of well-separated results should be very different from the original mixture. Therefore, the corresponding mSCM will be in a low position. It is noted that when the input mixture has a high input SNR, the lower mSCM constraint will filter out its separated results. Even though, the lower mSCM hypothesis still makes sense, because the filtered speech with high input SNR is somehow homogeneous and has limited benefits to model adaptation. In addition, the high input

SNR cases are rare for cross-domain task. Therefore, the lower mSCM constraint is safe and effective in most conditions.

### 3.2.2 Pseudo ground-truth selection

After computing both SCM and mSCM statistics of input mixtures, we re-organize all the statistics and speech signals that related to each unlabeled input mixture in a new data tuple format to facilitate the pseudo ground-truth selection. As shown in Fig. 2A, we call each data tuple as a “separation consistency information (SCI)” tuple, and it is organized as:

$$SCI = \{ID, SCM, mSCM, Mix, Sep1, Sep2\} \quad (4)$$

where ID is the mixture ID, Mix is the input mixture speech signal, Sep1 and Sep2 are the two individual speech signals that separated by DPCCN. With these SCI tuples, we then perform the pseudo ground-truth selection in two ways:

- CPS-1: Select SCI pairs with SCM value lies in the top  $p\%$  SCM range,  $p \in [0, 100]$ .
- CPS-2: Select SCI tuples with the following constraint:

$$SCI_s = \{SCI_k \mid (SCM_k > \alpha) \cap (mSCM_k < \beta)\} \quad (5)$$

where  $k = 1, 2, \dots, N$ ,  $SCI_s$  and  $SCI_k$  are the selected SCI tuples and  $k$ -th SCI tuple, respectively.  $\alpha, \beta$  are thresholds for SCM and mSCM, respectively.

No matter for CPS-1 or CPS-2, the separated signals, Sep1 and Sep2, in all the selected SCI tuples will be taken as the high confidence pseudo ground-truth for their corresponding mixture Mix. Then the selected mixtures with pseudo ground-truths are taken to form the “D-Pseudo Labeled Set” (pseudo ground-truth that produced by DPCCN) for further separation model adaptation. As discussed in the definition of mSCM, compared with CPS-1, perhaps CPS-2 is better at dealing with the difficult separation cases to some extent.

### 3.3 Heterogeneous Knowledge Fusion

The heterogeneous knowledge fusion (HKF), illustrated in Fig. 2D is used during the cross-knowledge adaptation in SCT-2 and SCT-3. HKF is a very simple operation just by replacing Sep1 and Sep2 in the selected SCI tuples of Fig. 2A with the outputs of the adapted Conv-TasNet as in SCT-2 and SCT-3. We use  $v_n^i$  to represent the  $i$ -th individual signal of  $n$ -th mixture separated by the adapted Conv-TasNet. The updated new data tuples  $\{Mix, Sep1, Sep2\}$  are then picked to form the “T-Pseudo Labeled Set” (pseudo ground-truths that

**Table 1** Cross-domain dataset information

Dataset	Acoustics	Type	#Spks	#Utts	Hours	Oracle
Libri2Mix (Source)	English/A	Train	251	13,900	58	✓
		Dev	40	3000	11	✓
		Test	40	3000	11	✓
WHAMR! (Target)	English/R	Train	101*	20,000	30	-
		Dev	101*	5000	10	-
		Test	18	3000	5	✓
Aishell2Mix (Target)	Mandarin/A	Train	340	10,000	11	-
		Dev	40	3000	3	-
		Test	20	3000	3	✓

"A" and "R" refer to anechoic and reverberant, respectively. "Oracle" indicates whether the oracle (ground-truth) data is available. "\*" Means the speakers of different sets are same

produced by Conv-TasNet). By doing so, the complementary information between the prior knowledge of separation consistency information that captured in the CPS block and the adapted Conv-TasNet are subtly integrated to further refine the primary DPCCN.

## 4 Experimental setups

### 4.1 Dataset

The publicly available English Libri2Mix [36] is used as our source domain dataset. Libri2Mix is a recent released anechoic separation corpus that contains artificial mixed speech from Librispeech [49]. We use the Libri2Mix generated from "train-100" subset to train our models. Two target domain datasets are used to validate our proposed methods, one is the English WHAMR! [37], the other is the Mandarin Aishell2Mix [35]. WHAMR! is a noisy and reverberant version of the WSJ0-2mix dataset [1] with four conditions (clean and anechoic, noisy and anechoic, clean and reverberant, noisy and reverberant). We take the clean and reverberant condition to evaluate the cross-domain speech separation performance. Note that the evaluation references of WHAMR! are also reverberant rather than anechoic. Aishell2Mix is created by ourselves [35], it is anechoic and released in [50]. Each mixture in Aishell2Mix is generated by mixing two-speaker utterances from Aishell-1 [51]. These utterances are randomly clamped to 4 seconds and rescaled to a random relative SNR between 0 and 5 dB. All datasets used in this study are resampled to 8kHz. The mixtures in both target domain datasets, WHAMR! and Aishell2Mix, are taken as the real-world unlabeled speech. Only the ground-truth of test sets in WHAMR! and Aishell2Mix are available for evaluating the speech separation performance, the training and development sets are all unlabeled. More details can be found in Table 1. It is worth noting that, the target domain development sets used to supervise the

model adaptation are also with pseudo ground-truth that produced by our proposed SCT.

### 4.2 Configurations

We keep the same network configurations of Conv-TasNet and DPCCN as in [13, 35], respectively. The model parameters of Conv-TasNet and DPCCN are 8.8M<sup>1</sup> and 6.3M. When processing a 4-s speech, the number of multiply-accumulate (MAC) operations [52] of Conv-TasNet and DPCCN are 28.2G and 33.1G, which are evaluated using open-source toolbox [53]. All models are trained with 100 epochs on 4-s speech segments. The initial learning rate is set to 0.001 and halved if the accuracy of development set is not improved in 3 consecutive epochs. Adam [54] is used as the optimizer and the early stopping is applied for 6 consecutive epochs. We use the standard negative SI-SNR [38] as loss function to train all separation systems. Utterance-level permutation invariant training (uPIT) [3] is used to address the source permutation problem. All source model adaptation related experiments are finished within 20 epochs. A Pytorch implementation of our DPCCN system can be found in [55].

### 4.3 Evaluation metrics

As our task is to improve cross-domain unsupervised speech separation, the performance improvement over the original mixture is more meaningful. Therefore, we report the well-known signal-to-distortion ratio improvement (SDRi) [56] and scale-invariant signal-to-noise ratio improvement (SI-SNRi) [38] to evaluate our proposed method.

<sup>1</sup> <https://arxiv.org/pdf/1809.07454v1.pdf>



**Table 2** SDRI/SI-SNRi (dB) performance of Conv-TasNet and DPCCN on Libri2Mix, WHAMR!, and Aishell2Mix test sets. Systems are all trained on Libri2Mix

System	SDRI/SI-SNRi (dB)		
	Libri2Mix	WHAMR!	Aishell2Mix
Conv-TasNet	12.41/11.98	6.83/6.45	2.57/2.08
DPCCN	13.48/13.04	8.99/8.50	5.78/5.09

## 5 Results and analysis

### 5.1 Cross-domain baselines

#### 5.1.1 Baselines

Both Conv-TasNet and DPCCN are taken as our cross-domain baseline systems. Performance is evaluated on all the in-domain Libri2Mix, and cross-domain WHAMR! and Aishell2Mix test sets. Results are shown in Table 2, where all separation systems are trained only on the Libri2Mix.

From Table 2, three findings are observed:

- 1) Compared with the performance on the in-domain Libri2Mix test set, there are huge cross-domain performance gaps exist on both the English and Mandarin target domain datasets.
- 2) Separation performance degradation caused by the language mismatch is much more severe than the acoustic reverberation.
- 3) DPCCN always shows much better speech separation performance than Conv-TasNet under both in-domain and cross-domain conditions.

The first two findings confirm that the current speech separation systems are very sensitive to cross-domain conditions, either for the time-domain Conv-TasNet, or the time-frequency domain DPCCN. The third observation shows the better system robustness of DPCCN over Conv-TasNet. We believe that the robustness gain of DPCCN mainly comes from using spectrogram to represent speech. For complicated tasks, such a handcrafted signal representation can provide more stable speech features than network learning. That's why we take the DPCCN individual outputs as references to calculate SCM for pseudo ground-truth selection as described in Section 3.2. We believe more reliable separation hypotheses can result in better pseudo ground-truth.

#### 5.1.2 Training with ground-truth labels

For results comparison and analysis, we also report the supervised separation performance of Conv-TasNet and DPCCN that trained with ground-truth labels in

**Table 3** Performance of Conv-TasNet and DPCCN trained with ground-truth labels on WHAMR! and Aishell2Mix test sets

System	Dataset	SDRI (dB)	SI-SNRi (dB)
Conv-TasNet	WHAMR!	11.03	10.59
	Aishell2Mix	9.00	8.32
DPCCN	WHAMR!	11.01	10.56
	Aishell2Mix	8.86	8.14

**Table 4** SDRI/SI-SNRi performance of DPCCN with CPS-1 and SCT-1 on Aishell2Mix test set

Top $p\%$ SCM	Adaptation	SDRI (dB)	SI-SNRi (dB)
$p = 10\%$	-	5.24	4.49
$p = 25\%$	-	5.65	5.05
$p = 50\%$	-	5.66	5.11
	✓	<b>5.98</b>	<b>5.32</b>

"-" means training model from scratch with only pseudo labeled data. "✓" means adapting model with the combined pseudo labeled data and the source domain Libri2Mix

Table 3, where all separation systems are trained with in-domain ground-truth sources of WHAMR! and Aishell2Mix. Interestingly, on the reverberant WHAMR! dataset, DPCCN and Conv-TasNet achieve almost the same results, while on the Aishell2Mix, DPCCN performs slightly worse than the Conv-TasNet. Coupled with the better cross-domain separation behaviors in Table 2, we take the DPCCN as our primary system, and the Conv-TasNet as the reviewer in all our following experiments.

### 5.2 Performance evaluation of SCT on Aishell2Mix

From Table 2 baseline results, we see the domain mismatch between English and Mandarin datasets is much larger than the two different English datasets. Therefore, in this section, we choose to first examine the proposed SCT on the Libri2Mix-Aishell2Mix (source-target) unsupervised cross-domain task, including evaluating the consistent pseudo-labeling and selection methods, CPS-1 and CPS-2, and different SCT variants for unsupervised model adaptation. Then, the optimized SCT is generalized to the WHAMR! dataset in Section 5.3.

#### 5.2.1 Initial examination of CPS-1

The DPCCN performance of the first unlabeled mixture pseudo label selection method, CPS-1, is first examined under SCT-1 framework in Table 4. Results of line 1-3 are from DPCCN that trained from scratch using CPS-1 outputs. These outputs are the "D-Pseudo Labeled Set" in SCT-1 with top  $p\%$ SCM target domain Aishell2Mix

**Table 5** SDRI/SI-SNRi (dB) performance of Conv-TasNet and DPCCN on Aishell2Mix test set under different SCT configurations

SCT	System	#Iter	CPS-1	CPS-2	Oracle
SCT-1	Conv-TasNet	1	5.14/4.63	5.47/4.90	5.98/5.39
		2	5.45/4.94	5.99/5.39	6.18/5.57
	DPCCN	1	5.98/5.32	5.90/5.25	6.00/5.31
		2	6.17/5.50	6.03/5.39	6.10/5.44
SCT-2	Conv-TasNet	1	5.14/4.63	5.47/4.90	5.98/5.39
		2	5.36/4.89	<b>6.15/5.52</b>	6.21/5.65
	DPCCN	1	6.05/5.52	<b>6.48/5.82</b>	6.79/6.19
		2	5.49/5.05	6.43/5.81	6.45/5.91
SCT-3	Conv-TasNet	1	5.14/4.63	5.47/4.90	-
		2	5.43/4.93	5.77/5.24	-
	DPCCN	1	6.14/5.58	6.22/5.65	-
		2	6.02/5.52	6.10/5.56	-

“Oracle” means using ground-truth as reference to calculate SI-SNR of separation outputs for selecting the pseudo ground-truth. All source models are well pre-trained on Libri2Mix. The best setup of  $\{\alpha, \beta\}$  in CPS-2 are  $\{5, 5\}$ ,  $\{8, 5\}$  in the 1st and 2nd iteration for all SCT variants, respectively.  $\eta$  is set to 5 for “Oracle selection”

data. We find that the separation performance can be improved by increasing the pseudo labeled training mixtures. And when  $p = 50\%$ , compared with the  $p = 25\%$  case, the additional performance improvements are rather limited even with an additional 25% data. Moreover, results of the last line show that, instead of training DPCCN from scratch, using the combined “D-Pseudo Labeled Set” and “Simulation Training Set” (Libri2Mix) to refine the source model (shown in Table 2, SDRI/SI-SNRi are 5.78/5.09, respectively) can further improve the target domain separation. In the following experiments, we set  $p = 50\%$  for all the CPS-1 experiments, and use Libri2Mix training set together with the “Pseudo Labeled Set” to fine-tune the source separation models for target model adaptation.

### 5.2.2 Evaluating SCT variants with both CPS-1 and CPS-2

Unlike only adapting DPCCN model as in the above CPS-1 initial experiments, in Table 5, we present the performance of both the updated target DPCCN and Conv-TasNet in each SCT iteration for all the three types of SCT variants. Experiments are still performed on the English-Mandarin cross-domain speech separation task. All source models are pre-trained on the same supervised Libri2Mix, then adapted to the Aishell2Mix condition using SCT-1 to SCT-3 frameworks separately. Besides the CPS-1 and CPS-2, in Table 5, we also report “oracle selection” performance using ground-truth as reference to calculate SI-SNR of separation outputs for selecting the pseudo ground-truth.

This “oracle selection” performance can be taken as the upper bound of our pseudo-labeling with heterogeneous neural network architecture. Two oracle selection criteria are used in our experiments: for SCT-1, we always calculate the best assignment SI-SNR between DPCCN outputs and ground-truth, while for SCT-2 and SCT-3, we use the SI-SNR scores between the ground-truth and DPCCN, Conv-TasNet outputs separately to select their corresponding individual separation signals as pseudo ground-truth, respectively. The pseudo ground-truth selection threshold  $\eta = 5$  is unchanged for each iteration in “oracle selection.” It is worth noting that, the  $\{\alpha, \beta, \eta\}$  are kept the same for both the pseudo-labeling of unlabeled training and development datasets.

From the English-Mandarin cross-domain separation results in Table 5, we can conclude the following observations:

- 1) Overall performance: Compared with baselines in Table 2, the best SCT variant, SCT-2 with CPS-2, improves the unsupervised cross-domain separation performance significantly. Specifically, absolute 3.68/3.44 dB and 0.70/0.73 dB SDRI/SI-SNRi improvements are obtained for Conv-TasNet and DPCCN, respectively. Moreover, the best performance of SCT-1 and SCT-2 with CPS-2 are very close to the upper bound ones with “oracle selection,” even both the training and development mixtures of target domain are taken as unlabeled ones. Such promising results indicate the effectiveness of our proposed SCT for improving the unsupervised cross-domain speech separation.
- 2) Model robustness: Under all SCT cases, the absolute performance gains achieved by the adapted Conv-TasNet are much bigger than the ones from the adapted DPCCN. However, the best DPCCN is always better than the best Conv-TasNet, this is possibly due to the better robustness or generalization ability of our previously proposed DPCCN in [35].
- 3) Pseudo label selection criterion: The CPS-2 performance is better than CPS-1 in almost all conditions, which tells us that introducing mSCM constraint is helpful to alleviate the pseudo ground-truth errors that brought by CPS-1. Combining both SCM and mSCM in CPS-2 can produce better high confidence pseudo labels.
- 4) Cross-knowledge adaptation: Together with CPS-2, the SCT-2 achieves better results over SCT-1, either for the best Conv-TasNet results or for the DPCCN ones. It proves the importance of cross-knowledge adaptation for leveraging the complementary information between heterogeneous models to target domain models.

**Table 6** SDRI/SI-SNRI(dB) performance on WHAMR! test set with SCT-2

SCT	System	#Iter	CPS-2	Oracle
SCT-2	Conv-TasNet	1	8.28/7.85	8.64/8.28
		2	8.48/8.06	8.68/8.27
	DPCCN	1	9.26/8.81	9.31/8.86
		2	8.84/8.40	8.95/8.52

"Oracle" and  $\eta$  have the same meaning as in Table 5. All source models are well pre-trained on Libri2Mix. The best setup of  $\{\alpha, \beta, \eta\}$  are  $\{8, 5, 8\}$  and  $\{12, 5, 8\}$  in the 1st and 2nd SCT iteration, respectively

- 5) Number of SCT iteration: For SCT-1, both Conv-TasNet and DPCCN are continuously improved in the first two iterations. For SCT-2 and SCT-3, Conv-TasNet still benefits from the 2nd iteration, while DPCCN only needs one iteration model adaptation to achieve the best results. This phenomenon indicates that the complementary cross-knowledge between different models can help DPCCN converge faster and achieve better performance.
- 6) Necessity of two-stage CPS-2: With CPS-2, SCT-3 does not bring any improvements over SCT-2, it means that the 2nd CPS-2 stage in SCT-3 is useless. Possibly because the updated Conv-TasNet has been refined by the first stage CPS-2 outputs, the new individual separation hypothesis of this updated model has homogeneous acoustic characteristic with the ones in the first stage CPS-2, resulting in relatively simple and partial separated pseudo ground-truth in the 2nd stage CPS-2. Considering this phenomenon, we stop trying more CPS stages and iterations in SCT pipelines, as feeding more homogeneous data is time-consuming and hard to bring additional benefits.

### 5.3 Performance evaluation of SCT on WHAMR!

As the SCT-2 with CPS-2 achieves the best results in Table 5, we generalize this framework to Libri2Mix-WHAMR! (source-target) task for a further investigation of unsupervised cross-domain speech separation. Both source and target domain are English speech mixtures but with different acoustic environments. Results are shown in Table 6. It's clear that we can obtain consistent observations from this table with the ones on Aishell2Mix, which verifies the good robustness and generalization ability of SCT under different cross-domain speech separation tasks. This nature of SCT is very important for real unsupervised speech separation applications. Our following experiments and analysis are all based on the

**Table 7** Overall SDRI/SI-SNRI(dB) performance with different configurations

Dataset	System	Baseline	SCT	Supervised
Aishell2Mix	Conv-TasNet	2.57/2.08	6.15/5.52	9.00/8.32
	DPCCN	5.78/5.09	6.48/5.82	8.86/8.14
WHAMR!	Conv-TasNet	6.83/6.45	8.48/8.06	11.03/10.59
	DPCCN	8.99/8.50	9.26/8.81	11.01/10.56

"Baseline" means model trained on source domain Libri2Mix while evaluated on target domain Aishell2Mix and WHAMR!. "SCT" is the best adaptation configuration, i.e. SCT-2 with CPS-2. "Supervised" means model trained with ground-truth labels

best SCT variant, SCT-2 with CPS-2, unless otherwise stated.

### 5.4 Overall performance evaluation

To better understand the proposed SCT, we re-organize the key experimental results in Table 7 for an overall comparison, including results of cross-domain baselines (in Table 2), the best SCT configuration (SCT-2 with CPS-2, in Tables 5 and 6), and the supervised results (upper bound) that trained with ground-truth labels (in Table 3). It is clear that the proposed SCT improves cross-domain separation performance significantly. Compared with Conv-TasNet, the SCT gain of DPCCN is much smaller. This may because the baseline performance of Conv-TasNet is much worse, when adapted with pseudo-labeled data pairs, Conv-TasNet will gain much more benefits. Besides, either for Conv-TasNet or DPCCN, the selected data during SCT actually has similar acoustic characteristics. This means that after SCT adaptation, the target domain performance of Conv-TasNet and DPCCN would reach to a similar level (as shown in the SCT column). In addition, results in Table 7 indicate that there is still a big performance gap between SCT and the upper bound ones, which motivates us to further improve the current SCT in our future works. Even though, considering the huge performance gain of SCT over baseline, we still believe the SCT is promising for tackling unsupervised speech separation tasks.

### 5.5 Heterogeneous separation results fusion

Motivated by the design of SCT, we believe that the separation results of the final adapted target domain models also have complementary information, because they are derived from two different neural networks with heterogeneous structure. Therefore, a simple linear fusion of separated signal spectrograms is preliminarily investigated to further improve the SCT.

Results are shown in Table 8, where  $\lambda$  and  $1 - \lambda$  are linear weights for the signal spectrograms of adapted DPCCN and Conv-TasNet outputs respectively. The

**Table 8** Performance of heterogeneous separation results fusion on Aishell2Mix and WHAMR! test sets

Dataset	$\lambda$	SDRi (dB)	SI-SNRi (dB)
Aishell2Mix	0.5	6.55	5.93
	0.8	<b>6.57</b>	<b>5.93</b>
	0.9	6.53	5.88
WHAMR!	0.5	9.19	8.79
	0.8	<b>9.34</b>	<b>8.92</b>
	0.9	9.32	8.89


pairwise cosine similarity is used to find the best match spectrograms that belong to the same speaker during linear fusion. Compared with the best SCT-2 results in Tables 5 and 6, this simple fusion is still able to bring slight performance improvements. This indicates that, it is possible to exploit the complementary information between SCT outputs to further improve the final separation results. It will be interesting to try other and more effective separation results fusion methods in future works.

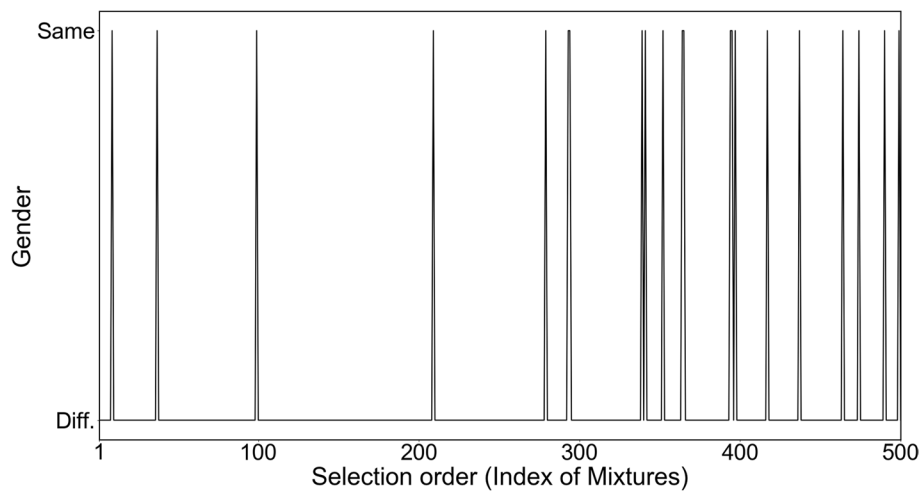
### 5.6 Data quantity analysis of pseudo ground-truth

The key success of the proposed SCT depends on the high confidence pseudo-labeling. It's very important to

analyze the data amount statistics of the selected pseudo ground-truth during SCT in different unsupervised separation tasks. Figure 3 shows the statistics that used to adapt the heterogeneous networks during each iteration of SCT-2 (with CPS-2) in Tables 5 and 6, including the selected training and development data of unlabeled Aishell2Mix and WHAMR! datasets. For further comparisons, we also show the corresponding upper bound data statistics generated by the “Oracle selection” as references. Note that, as the cross-knowledge adaptation is applied during SCT-2, the data amounts of “D-Pseudo Labeled Set” and “T-Pseudo Labeled Set” are the same but with different ground-truth individual signals, so we use “SCT-2” to represent both of them, and the “Oracle Conv-TasNet” and “Oracle DPCCN” in Fig. 3 actually represent the oracle amount of pseudo data that selected to adapt the Conv-TasNet and DPCCN, respectively.

From Fig. 3, three findings are observed: (1) the 2nd SCT-2 iteration can produce more high confidence data, and the selected data quantity is close to the upper bounds with “Oracle selection,” indicating the heterogeneous structure in SCT and the thresholds of CPS-2 are reasonable; (2) on Aishell2Mix, both the selected training and development data increments in the 2nd iteration are higher than the ones on WHAMR!, which means the

**Fig. 3** Data quantity of selected pseudo ground-truth of SCT-2 (with CPS-2) versus the “Oracle selection” on Aishell2Mix and WHAMR! unlabeled training and development sets



**Fig. 4** Gender information of top 500 CPS-2 results on Aishell2Mix development set. “Diff.” represents the different gender


multiple SCT-2 iterations are necessary for tasks with the larger cross-domain mismatch. (3) for “Oracle DPCCN,” the selected data quantities of two iterations are almost the same, indicating the pseudo-labeled mixtures in each iteration are a large number of homogeneous data that results in an over-trained DPCCN model. This is also the reason of worse results in the 2nd iteration that shown in Tables 5 and 6. All these above findings give a well support to the separation results as presented in both Tables 5 and 6.

### 5.7 Gender preference analysis

As we all know, the speech mixed with different gender speakers is easier to separate than that with the same


gender speakers. In this section, we investigate the gender distribution of selected pseudo-labels on the Aishell-2Mix development set. The gender information of top 500 mixtures with the best CPS-2 setup,  $\alpha = 8$  and  $\beta = 5$ , is presented in Fig. 4, where each spike pulse represents the gender in each mixture changing from different to the same.

From Fig. 4, it is clear that the proposed CPS-2 prefers to select the mixtures with different gender speakers. The sparse spike pulse shows the extremely low proportion of same gender mixtures in the entire selected speech, and its distribution tends to denser when the confidence of the selected mixture becomes lower (larger selection order). These phenomena are consistent with our prior



**Fig. 5** SI-SNRi (dB) of DPCCN and Conv-TasNet separation results of the 310 “bad cases” varies with the separation consistency measure





**Fig. 6** Detailed separation statistics of the 310 “bad cases” in Aishell2Mix development set

knowledge, i.e., the speech mixed by different gender speakers is easier to separate and its separated signals from heterogeneous models show a higher separation consistency.

### 5.8 Bad cases analysis

Finally, we perform a bad cases analysis of the separation results on the Aishell2Mix development set in Fig. 5. All these unlabeled mixtures in this dataset are first separated by the best adapted target domain DPCCN and Conv-TasNet models in Table 5 (SCT-2 with CPS-2). Then the CPS-2 with  $\alpha = 8$ ,  $\beta = 5$  is used to select the pseudo labels and 1716 mixtures’ SCI tuples are selected in total. Next, we calculate the standard separation performance (SI-SNRi) of both the DPCCN and Conv-TasNet separation outputs by taking the real ground-truth to evaluate each mixture performance, and we refer them to  $\text{SI-SNRi}_{\text{DPCCN}}$  and  $\text{SI-SNRi}_{\text{Conv-TasNet}}$  for simplicity. Then, we compare each SI-SNRi with the average SI-SNRi (5.52 dB, the best performance of Conv-TasNet in Table 5) of Aishell2Mix test set to determine whether the current mixture separation is a “bad case” or not. For each selected mixture, if its  $\{\text{SI-SNRi}_{\text{DPCCN}} \parallel \text{SI-SNRi}_{\text{Conv-TasNet}}\} < 5.52$  dB, we consider it a failure separation (F) and the corresponding mixed speech is taken as a “bad case,” otherwise we take it as a success separation (T). With this rule, total 310 of 1716 (18.1%) mixtures are taken as “bad cases.”

The reason behind this “bad case” decision rule is that, in the speech separation field, there is no measurement to evaluate each speech separation is 100% accurate or not. Therefore, we think that, the real separation performance of the best separation model can be taken as a proper heuristic signal distortion threshold for a rough “bad case” analysis. And in our SCT-2, when compared with the best DPCCN performance (5.82 dB) in Table 5, the Conv-TasNet performance, 5.52 dB is a stricter one for the “bad case” decision.

Figure 5 shows how the DPCCN and Conv-TasNet separation outputs of the 310 “bad cases” SI-SNRi varies with the separation consistency SCM. From these scatter points, we see that, with our proposed CPS-2, the selected 310 mixture pseudo labels still contain low-quality ones that are not suitable to be taken as ground-truth, even though all these mixtures have relatively high consistency confidence. From the left part of this figure, we find some “bad cases” with high separation consistency  $\text{SCM} > 12$  dB but their real separation performances are very low ( $\text{SI-SNRi} < 2$  dB). However, on the contrary, the right part figure shows some low SCM mixtures are also separated very well. Therefore, we speculate that, these “bad cases” may not be too bad if they are within the error tolerance of system training data, they may be taken as small noisy distortions of the whole pseudo labeled training set and may help to enhance the model robustness. That’s why we still obtain promising performance in Table 5 using the proposed SCT.

Figure 6 demonstrates other detailed separation statistics of the same 310 “bad cases” on Aishell2Mix development set from another perspective. The T, F means the success, failure separation as defined in the above statements. Each “bad case” covers three kinds of T, F combination, such as,  $\text{Conv-TasNet}(T) \cap \text{DPCCN}(F)$  means for each unlabeled mixture, the separation of Conv-TasNet is success while DPCCN is failure.

From Fig. 6, we see 56.8% of these “bad cases” are consistent failure separations for both DPCCN and Conv-TasNet. However, there is still around half of the data can be separated well by one of these two heterogeneous systems, as shown in the two T∩F combinations. This observation clearly proves the large complementary information between two heterogeneous separation models, as the time-domain Conv-TasNet and the time-frequency domain DPCCN used in our SCT. And it also inspires us to improve the SCT-1 to SCT-2 using the cross-knowledge adaptation. Besides, for the 31.3% vs 11.9% T∩F combination, we see there are much more

DPCCN success mixture separations than the Conv-TasNet on this difficult-to-separate 310 mixtures. This means DPCCN is a better candidate for robust speech separation task, using DPCCN as the primary model and its outputs as references in the whole SCT process is reasonable.

## 6 Conclusion

In this paper, we proposed an iterative separation consistency training (SCT) framework, a practical source model adaptation technology for cross-domain unsupervised speech separation tasks. By introducing an effective pseudo-labeling approach, the unlabeled target domain mixtures are well exploited for target model adaptation, which successfully reduces the strong ground-truth reliance of most state-of-the-art supervised speech separation systems.

Different from previous works, SCT follows a heterogeneous structure, it is composed of a masking-based time-domain separation model, Conv-TasNet, and a mapping-based time-frequency domain separation model, DPCCN. Due to this heterogeneous structure and the specially designed separation consistency measures, SCT can not only perform the pseudo-labeling of unlabeled mixtures automatically, but also can ensure the selected pseudo ground-truths are high quality and informative. Moreover, by introducing the cross-knowledge adaptation in SCT, the large complementary information between heterogeneous models is maximally leveraged to improve the primary separation system. In addition, the iterative adaptation nature in SCT provides an increased chance to improve the primary model when there is a large amount of unlabeled mixtures available. Finally, we find this heterogeneous design of SCT also has the potential to further improve the final separation system performance by combing two final adapted separation model at the level of their outputs.

We verified the effectiveness of our proposed methods on two cross-domain conditions: the reverberant English and the anechoic Mandarin acoustic environments. Results show that, under each condition, both the heterogeneous separation models are significantly improved, their target domain performance is very close to the upper bound ones, even the target domain training and development sets are all unlabeled mixtures. In addition, through the bad case analysis, we find that the SCT will definitely introduce some error pseudo ground-truth to a certain extent. This limitation of current SCT still needs to be improved in our future works before we apply it to real speech separation applications.

## Acknowledgements

Not applicable.

## Authors' contributions

Jiangyu implemented the SCT methods and wrote the manuscript, which then was revised by yanhua. The whole project conducted under yanhua's supervision. They were both major contributors of this work. All authors read and approved the manuscript.

## Funding

This work is supported by the National Natural Science Foundation of China (Grant No.62071302 and No.61701306).

## Availability of data and materials

The three datasets used in this article are publicly archived datasets: <https://github.com/JorisCos/LibriMix>; <https://wham.whisper.ai>; <https://github.com/jyhan03/icassp22-dataset>.

Some visualization materials can be found in: <https://github.com/jyhan03/icassp22-dataset/tree/main/sct>.

## Declarations

### Competing interests

The authors declare that they have no competing interests.

Received: 30 October 2022 Accepted: 12 January 2023

Published online: 20 January 2023

## References

1. J.R. Hershey, Z. Chen, J. Le Roux, S. Watanabe, in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Deep clustering: Discriminative embeddings for segmentation and separation (IEEE, 2016), pp. 31–35
2. D. Yu, M. Kolbæk, Z.H. Tan, J. Jensen, in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Permutation invariant training of deep models for speaker-independent multi-talker speech separation (IEEE, 2017), pp. 241–245
3. M. Kolbæk, D. Yu, Z.H. Tan, J. Jensen, Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks. *IEEE/ACM Trans. Audio Speech Lang. Process.* **25**(10), 1901–1913 (2017)
4. Z. Chen, Y. Luo, N. Mesgarani, in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Deep attractor network for single-microphone speaker separation (IEEE, 2017), pp. 246–250
5. Y. Luo, Z. Chen, N. Mesgarani, Speaker-independent speech separation with deep attractor network. *IEEE/ACM Trans. Audio Speech Lang. Process.* **26**(4), 787–796 (2018)
6. Z.Q. Wang, J. Le Roux, J.R. Hershey, in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Alternative objective functions for deep clustering (IEEE, 2018), pp. 686–690
7. D. Wang, J. Chen, Supervised speech separation based on deep learning: An overview. *IEEE/ACM Trans. Audio Speech Lang. Process.* **26**(10), 1702–1726 (2018)
8. D.S. Williamson, Y. Wang, D. Wang, Complex ratio masking for monaural speech separation. *IEEE/ACM Trans. Audio Speech Lang. Process.* **24**(3), 483–492 (2015)
9. K. Tan, D. Wang, in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Complex spectral mapping with a convolutional recurrent network for monaural speech enhancement (IEEE, 2019), pp. 6865–6869
10. H.S. Choi, J.H. Kim, J. Huh, A. Kim, J.W. Ha, K. Lee, Phase-aware speech enhancement with deep complex u-net. *arXiv preprint arXiv:1903.03107* (2019)
11. Y. Hu, Y. Liu, S. Lv, M. Xing, S. Zhang, Y. Fu, J. Wu, B. Zhang, L. Xie, in *Proc. Interspeech*. DCCRN: Deep complex convolution recurrent network for phase-aware speech enhancement (IEEE, 2020), pp. 2472–2476
12. Y. Fu, J. Wu, Y. Hu, M. Xing, L. Xie, in *2021 IEEE Spoken Language Technology Workshop (SLT)*. DESNet: A multi-channel network for simultaneous speech dereverberation, enhancement and separation (IEEE, 2021), pp. 857–864

13. Y. Luo, N. Mesgarani, Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation. *IEEE/ACM Trans. Audio Speech Lang. Process.* **27**(8), 1256–1266 (2019)
14. Y. Luo, Z. Chen, T. Yoshioka, in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Dual-path RNN: efficient long sequence modeling for time-domain single-channel speech separation (IEEE, 2020), pp. 46–50
15. J. Zhang, C. Zorilă, R. Doddipatla, J. Barker, in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. On end-to-end multi-channel time domain speech separation in reverberant environments (IEEE, 2020), pp. 6389–6393
16. J. Chen, Q. Mao, D. Liu, in *Proc. Interspeech*. Dual-path transformer network: Direct context-aware modeling for end-to-end monaural speech separation, pp. 2642–2646 (ISCA, Shanghai, China, 2020)
17. C. Subakan, M. Ravanelli, S. Cornell, M. Bronzi, J. Zhong, in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Attention is all you need in speech separation (IEEE, 2021), pp. 21–25
18. J. Han, X. Zhou, Y. Long, Y. Li, in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Multi-channel target speech extraction with channel decorrelation and target speaker adaptation (IEEE, 2021), pp. 6094–6098
19. J. Han, W. Rao, Y. Wang, Y. Long, in *Proc. Interspeech*. Improving channel decorrelation for multi-channel target speech extraction, pp. 1847–1851 (ISCA, Brno, Czech, 2021)
20. J. Han, W. Rao, Y. Long, J. Liang, in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. Attention-based scaling adaptation for target speech extraction (IEEE, 2021), pp. 658–662
21. J. Chen, Q. Mao, D. Liu, in *Proc. Interspeech*. On synthesis for supervised monaural speech separation in time domain, pp. 2627–2631 (ISCA, Shanghai, China, 2020)
22. Z.Q. Wang, D. Wang, in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Count and separate: Incorporating speaker counting for continuous speaker separation (IEEE, 2021), pp. 11–15
23. Y. Luo, Z. Chen, C. Han, C. Li, T. Zhou, N. Mesgarani, in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Rethinking the separation layers in speech separation networks (IEEE, 2021), pp. 1–5
24. G.P. Yang, C.I. Tuan, H.Y. Lee, L.S. Lee, in *Proc. Interspeech*. Improved speech separation with time-and-frequency cross-domain joint embedding and clustering, pp. 1363–1367 (ISCA, Graz, Austria, 2019)
25. X.L. Zhang, D. Wang, A deep ensemble learning method for monaural speech separation. *IEEE/ACM Trans. Audio Speech Lang. Process.* **24**(5), 967–977 (2016)
26. H. Shi, L. Wang, S. Li, C. Ding, M. Ge, N. Li, J. Dang, H. Seki, in *Proc. Interspeech*. Singing voice extraction with attention-based spectrograms fusion, pp. 2412–2416 (ISCA, Shanghai, China, 2020)
27. H. Shi, L. Wang, M. Ge, S. Li, J. Dang, in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Spectrograms fusion with minimum difference masks estimation for monaural speech dereverberation (IEEE, 2020), pp. 7544–7548
28. S. Wisdom, E. Tzinis, H. Erdogan, R.J. Weiss, K. Wilson, J.R. Hershey, in *Advances in Neural Information Processing Systems*. Unsupervised sound separation using mixture invariant training (NeurIPS, Vancouver, Canada, 2020)
29. S. Wisdom, A. Jansen, R.J. Weiss, H. Erdogan, J.R. Hershey, in *Workshop on Applications of Signal Processing to Audio and Acoustics*. Sparse, efficient, and semantic mixture invariant training: Taming the-wild unsupervised sound separation (IEEE, New Paltz, USA, 2021)
30. A. Sivaraman, S. Wisdom, H. Erdogan, J.R. Hershey, Adapting speech separation to real-world meetings using mixture invariant training. *arXiv preprint arXiv:2110.10739* (2021)
31. J. Zhang, C. Zorilă, R. Doddipatla, J. Barker, in *Proc. Interspeech*. Teacher-student MixIT for unsupervised and semi-supervised speech separation, pp. 3495–3499 (ISCA, Brno, Czech, 2021)
32. R. Gao, K. Grauman, in *Proceedings of the IEEE/CVF International Conference on Computer Vision*. Co-separating sounds of visual objects, pp. 3879–3888 (IEEE, Seoul, Korea, 2019)
33. Y. Hoshen, in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Towards unsupervised single-channel blind source separation using adversarial pair unmixed-and-remix (IEEE, 2019), pp. 3272–3276
34. M.W. Lam, J. Wang, D. Su, D. Yu, in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Mixup-breakdown: a consistency training method for improving generalization of speech separation models (IEEE, 2020), pp. 6374–6378
35. J. Han, Y. Long, L. Burge, J. Cernocky, in *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. DPCCN: Densely-connected pyramid complex convolutional network for robust speech separation and extraction, pp. 7292–7296
36. J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, E. Vincent, Librimix: An open-source dataset for generalizable speech separation. *arXiv preprint arXiv:2005.11262* (2020)
37. M. Maciejewski, G. Wichern, E. McQuinn, J. Le Roux, in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. WHAMR!: Noisy and reverberant single-channel speech separation (IEEE, 2020), pp. 696–700
38. J. Le Roux, S. Wisdom, H. Erdogan, J.R. Hershey, in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Sdr-half-baked or well done? (IEEE, 2019), pp. 626–630
39. J. Heitkaemper, D. Jakobeit, C. Boeddeker, L. Drude, R. Haeb-Umbach, in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Demystifying TasNet: A dissecting approach (IEEE, 2020), pp. 6359–6363
40. S. Sonning, C. Schüldt, H. Erdogan, S. Wisdom, in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Performance study of a convolutional time-domain audio separation network for real-time speech denoising (IEEE, 2020), pp. 831–835
41. T. Ochiai, M. Delcroix, R. Ikeshita, K. Kinoshita, T. Nakatani, S. Araki, in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. Beam-TasNet: Time-domain audio separation network meets frequency-domain beamformer (IEEE, 2020), pp. 6384–6388
42. T. von Neumann, K. Kinoshita, L. Drude, C. Boeddeker, M. Delcroix, T. Nakatani, R. Haeb-Umbach, in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. End-to-end training of time domain audio separation and recognition (IEEE, 2020), pp. 7004–7008
43. C. Lea, R. Vidal, A. Reiter, G.D. Hager, in *European Conference on Computer Vision*. Temporal convolutional networks: A unified approach to action segmentation (Springer, 2016), pp. 47–54
44. C. Lea, M.D. Flynn, R. Vidal, A. Reiter, G.D. Hager, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Temporal convolutional networks for action segmentation and detection, pp. 156–165 (IEEE, Hawaii, USA, 2017)
45. S. Bai, J.Z. Kolter, V. Koltun, An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271* (2018)
46. O. Ronneberger, P. Fischer, T. Brox, in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. U-net: Convolutional networks for biomedical image segmentation (Springer, 2015), pp. 234–241
47. G. Huang, Z. Liu, L. Van Der Maaten, K.Q. Weinberger, in *Proceedings of the IEEE conference on computer vision and pattern recognition*. Densely connected convolutional networks, pp. 4700–4708 (IEEE, Hawaii, USA, 2017)
48. H. Zhao, J. Shi, X. Qi, X. Wang, J. Jia, in *Proceedings of the IEEE conference on computer vision and pattern recognition*. Pyramid scene parsing network, pp. 2881–2890 (IEEE, Hawaii, USA, 2017)
49. V. Panayotov, G. Chen, D. Povey, S. Khudanpur, in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. Librispeech: an asr corpus based on public domain audio books (IEEE, 2015), pp. 5206–5210
50. Aishell2Mix. <https://github.com/jyhan03/icassp22-dataset>. Accessed 18 Jan 2023
51. H. Bu, J. Du, X. Na, B. Wu, H. Zheng, in *2017 20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*. AISHELL-1: An open-source mandarin speech corpus and a speech recognition baseline (IEEE, 2017), pp. 1–5

52. N. Whitehead, A. Fit-Florea, Precision & performance: Floating point and IEEE 754 compliance for nvidia gpus. *rn (A+ B)* **21**(1), 18749–19424 (2011)
53. THOP. <https://github.com/Lyken17/pytorch-OpCounter>. Accessed 18 Jan 2023
54. D.P. Kingma, J. Ba, Adam: A method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
55. DPCCN. <https://github.com/jyhan03/dpccn>. Accessed 18 Jan 2023
56. E. Vincent, R. Gribonval, C. Févotte, Performance measurement in blind audio source separation. *IEEE Trans. Audio Speech Lang. Process.* **14**(4), 1462–1469 (2006)

## Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Submit your manuscript to a SpringerOpen<sup>®</sup> journal and benefit from:**

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

---

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)

---