

Show from Tell

Jiao, Jianbo; Alsharid, Mohammad; Drukker, Lior; Papageorghiou, Aris T.; Zisserman, Andrew; Noble, J. Alison

DOI:

[10.48550/arXiv.2310.16477](https://doi.org/10.48550/arXiv.2310.16477)

License:

None: All rights reserved

Citation for published version (Harvard):

Jiao, J, Alsharid, M, Drukker, L, Papageorghiou, AT, Zisserman, A & Noble, JA 2023 'Show from Tell: Audio-Visual Modelling in Clinical Settings' arXiv, pp. 1-12. <https://doi.org/10.48550/arXiv.2310.16477>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Show from Tell: Audio-Visual Modelling in Clinical Settings

Jianbo Jiao^{†*,1,2}, Mohammad Alsharid^{†,1,3}, Lior Drukker^{4,5}, Aris T. Papageorgiou⁴, Andrew Zisserman¹, and J. Alison Noble^{*,1}

¹Department of Engineering Science, University of Oxford, Oxford, UK

²School of Computer Science, University of Birmingham, Birmingham, UK

³Department of Electrical Engineering and Computer Science, Khalifa University, Abu Dhabi, United Arab Emirates

⁴Nuffield Department of Women's & Reproductive Health, University of Oxford, Oxford, UK

⁵Rabin Medical Center, Tel-Aviv University Faculty of Medicine, Israel

ABSTRACT

Auditory and visual signals usually present together and correlate with each other, not only in natural environments but also in clinical settings. However, the audio-visual modelling in the latter case can be more challenging, due to the different sources of audio/video signals and the noise (both signal-level and semantic-level) in auditory signals – usually speech. In this paper, we consider audio-visual modelling in a clinical setting, providing a solution to learn medical representations that benefit various clinical tasks, without human expert annotation. A simple yet effective multi-modal self-supervised learning framework is proposed for this purpose. The proposed approach is able to localise anatomical regions of interest during ultrasound imaging, with only speech audio as a reference. Experimental evaluations on a large-scale clinical multi-modal ultrasound video dataset show that the proposed self-supervised method learns good transferable anatomical representations that boost the performance of automated downstream clinical tasks, even outperforming fully-supervised solutions.

Introduction

Visual signals are often accompanied by auditory signals which provide essential assistance for scene understanding. Such multi-modal scenarios widely exist in both natural environments and clinical settings. Our goal is to be able to learn visual capabilities (such as localising *kidneys*, or detecting standard planes in an ultrasound scan) in clinical scenarios simply by listening to sonographers (*i.e.* operators) describing what they are doing.

Audio-visual modelling has attracted interest in the recent literature [2, 3, 13], in which the basic assumption is that the underlying video and audio signals are densely correlated with each other. Based on this assumption, correspondence modelling is leveraged, *i.e.* examining if the input audio and video signals correspond or not. Such a binary discrimination model is simple, but has been shown to be effective in learning audio-visual representations. Some work [14, 17] has proposed contrastive learning for audio-visual modelling by referring to negative samples. This helps model discrimination ability and has been shown to result in a better representation quality [17].

Promising progress has been made in the above-mentioned literature for audio-visual modelling in natural scenes. However, in clinical scenarios the setting is different and consequently it is more challenging to model (see Figure 1a). Different from the natural scenario, in a clinical setting the auditory signal and visual signal usually arise from different sources, and have sparse correlation, *e.g.* the conversational speech audio and the corresponding captured medical images/videos. This prevents the direct application of existing methods that were designed for natural scenes. Similar to natural scenes, the modelling of audio-visual signals is helpful - in this case for learning medical representations without human supervision, especially expert manual annotations, which are costly and hard to scale. During real-time ultrasound (US) imaging narrations from clinicians are describing the visual signals and as a result they have underlying strong correspondence with the anatomy seen on screen. Our hypothesis is that modelling such correlations in a self-supervised manner will be beneficial for automating downstream clinical tasks.

To this end, in this paper, we consider audio-visual modelling in a clinical setting and propose a simple framework (Figure 1b) to address the challenges faced in doing this. Here the *clinical setting* we used is fetal US scanning, where audio recordings containing the speech uttered by the sonographers during the scan are considered, together with the US video shown on the screen. We can do this because sonographers speak about what they see on the screen during an examination, as an

[†]These authors contributed equally to this work

*Corresponding authors: {jianbo.jiao, alison.noble}@eng.ox.ac.uk




Figure 1. Multi-modal representation learning in clinical setting. (a) Illustration of audio-visual modelling in a natural scene setting (top) and in a clinical setting (bottom). (b) Pipeline for audio-video modelling and analysis in clinical settings from raw video footage. Video frames with the corresponding speech audio are extracted from the raw footage (I). After data pre-processing and text data generation, the enhanced multi-modal data (II) are fed into a joint fusion framework to learn multi-modal representations without human annotation (III). The whole automated system can be transferred to several downstream tasks and used for large-scale analysis (IV). (c) Evaluation results on the downstream task of standard plane detection. Note the method with dotted borders is fully-supervised using external human annotations.

explanation to the pregnant individual being scanned. More specifically, we start with modelling the first-order correspondence between the speech audio and its corresponding video signal. Cross-modal contrastive learning is then introduced to strengthen the discrimination between the current sample and unrelated negative samples. Clinical audio recordings are noisy. Therefore, we propose to pre-process the audio data to enhance the related speech content before feeding it to the deep model. To enrich the semantic understanding of the audio data and bootstrap the audio-visual modelling, a text branch is introduced to the proposed framework, forming a multi-modal modelling framework. The text is automatically transcribed from the audio signal, and we propose a learnable gating module towards the end of this text branch to spot the text embeddings of the relevant anatomies. Note that in our clinical ultrasound setting, the originally captured data are the fetal ultrasound video from the ultrasound machine and the speech audio from the sonographers. The text information is not available and accurate transcription is infeasible to achieve. Thus our work aims at audio-visual modelling and takes textual data as intermediate information to help the modelling.

In addition to learning self-supervised cross-modal representations, we demonstrate audio-guided visual localisation by directly leveraging the learned representations. This could be used in various clinical applications, *e.g.* better user experience for patients during scanning, training new sonographers and low-cost data annotation for region-of-interest by only using audio, to name a few. Note that during the inference stage (*i.e.* in real applications), no text information is needed and only the audio signal is sufficient to help the visual localisation. We perform extensive experimental evaluations on a large-scale multi-modal clinical US dataset to validate the effectiveness of the proposed approach by transferring the learned representations to downstream tasks. Analysis of three downstream clinical tasks shows that the proposed approach learns good transferable representations, which surpass the performance of fully-supervised solutions. To summarise, the main contributions of this work are as follows:

- We propose, to our knowledge, the first audio-visual modelling framework with generated text as an additional modality for representation learning, in a clinical setting where the correlation modelling for audio-visual data is more challenging than natural scenes.
- We introduce several strategies to address the challenges in clinical settings, including audio pre-processing, selective gating module for text embedding refinement and multi-modal joint optimisation.




Figure 2. Audio-video modelling architectures. (a) Baseline model of correspondence modelling. (b) Correspondence modelling with contrastive learning. (c) The proposed multi-modal modelling approach with additional textual information and visual-spatial awareness. (d-I) Basic fusion strategy by concatenating multi-modal features. (d-II) Spatial-aware fusion by element-wise product. (e-I) The proposed selective information gating (SIG) module, with the coarse level of only filtering outliers. (e-II) The SIG module with fine-grained level of keywords spotting by only focusing on specific keywords of interest. The \otimes means element-wise product.

- Experimental analysis validates the effectiveness of the proposed framework, with better performance than alternative solutions and superior to fully-supervised settings. We also present an approach to visually localise anatomical regions of interest according to audio guidance.

Methods

We propose a framework to model the audio-visual representations by leveraging intermediate text information. The proposed approach has three original components: audio concentration for audio data processing, a selective gating module for embedding refinement, and a multi-modal correlation optimisation loss. Figure 2 illustrates the framework of audio-video modelling with the main ideas of the proposed method shown in Figure 2c.

Large-Scale Multi-Modal Clinical Dataset

The proposed approach is validated on a large-scale multi-modal clinical US dataset called the PULSE (Perception Ultrasound by Learning Sonographer Experience) dataset which is described in [9]. Ethics approval was granted by the West of Scotland Research Ethics Service, UK Research Ethics Committee (Reference 18/WS/0051). Informed consent was obtained from all participants, and all methods were performed in accordance with the relevant guidelines and regulations. Briefly, the PULSE dataset consists of full-length US video recorded with concurrent speech audio, eye-gaze tracking data and probe motion data, as sonographers perform routine scans on pregnant women. A typical scan setting is shown in Figure 1a bottom. In this work, we focus on modelling the video and audio data for the second-trimester data (*i.e.* the 20-week anomaly scan). The video data was frame-grabbed from the US machine in real-time to a computer in a fully anonymised fashion. The original full-length video was stored at 30 FPS using lossless compression. The speech audio data was recorded in synchronisation with the video data by using microphones placed in the scanning room.

Multi-Modal Medical Representation Learning

Baselines We first present a baseline to model the correlation between the US video and audio data modalities. Specifically, the most intuitive approach is to predict if the given video and audio samples correspond or not, as usually done in natural scenes. To this end, we use a binary classification framework to model audio-visual correspondence. As illustrated in Figure 2a, the US video clip and audio segment are input to the video and audio sub-networks to obtain corresponding feature embeddings, followed by a fusion operation to produce the final feature for correspondence prediction. The original recorded video with the accompanying synchronised audio signals are assumed to correspond with each other, while shifted video-audio data pairs

or pairs from different scans are considered to be non-corresponding. Let the video signal be V and the audio signal A . The correspondence modelling task can be simply modelled as a binary classification problem:

$$L_{base}(V_i, A_i) = -\frac{1}{N} \sum_{i=1}^N c_i \cdot \log(p(A_i, V_i)) + (1 - c_i) \cdot \log(1 - p(A_i, V_i)), \quad (1)$$

where c is the correspondence label (0 or 1) and p is the predicted probability of the audio-video pair being in correspondence. N is the number of samples. Note here the label c does not need any human annotations, but instead naturally encoded in the data. Specifically, when the audio and video signals are aligned with each other (by default naturally as they are in the dataset), $c = 1$, otherwise, $c = 0$ if they have not been matched with the signal that they originally correspond to in the dataset.

Cross-modal contrastive learning The above baseline model only describes the alignment of the audio and video data, being a very strict constraint that may not be sufficient to model the relationship between these two modalities due to the challenges in the clinical setting. We observe that in our speech audio data, there are some medical-unrelated contents (*e.g.* random talk), which deteriorate the correlation learning between the two modalities. Contrastive learning aims to pull the similar data pairs closer while pushing dissimilar ones away from each other, in the embedding space during optimisation. In addition to modelling the relationship of the positive samples that are aligned at the same timestamp, it also considers dissimilar sample pairs, which act as negative samples for the modelling. As a result, built upon the baseline model, cross-modal contrastive learning is introduced to model the audio-video representation, and encourage a soft constraint towards the audio-visual relationship modelling. Specifically, consider a positive pair of audio-video data as (A_i, V_i) , and other sample pairs (A_i, V_j) that are misaligned (*i.e.* $i \neq j$) as the negative pairs. Then the loss function for contrastive learning L_{contra} is defined by:

$$L_{contra}(V_i, A_i) = -\log \frac{\text{sim}(V_i, A_i)}{\sum_{j=1}^N \mathbb{1}_{[j \neq i]} \text{sim}(V_j, A_i)}, \quad (2)$$

where $\mathbb{1}_{[j \neq i]}$ equals 1 iff $j \neq i$ and otherwise 0. The $\text{sim}()$ function is a similarity measurement over the exponential of the arguments. Though intra-modal contrastive learning may be able to boost the representation learning ability if further included, here we focus on cross-modal contrastive learning. And we found in a recent work [10] that it may lead to some issues in clinical settings, where additional anatomy annotations may be required. As a result, in this paper we focus on the cross-modal learning approach.

Spatial awareness Although contrastive learning provides an adaptive correlation modelling solution, the conventional similarity measure, *i.e.* a one-dimensional feature vector (as shown in Figures 2a and 2b), ignores spatial context in the feature embedding and loses local spatial awareness. To address this, instead of compressing the output bottleneck feature of the video network to be a 1-D vector, we preserve the spatial dimensions of the feature before fusing it with the audio feature (Figure 2c). Specifically, the spatial dimensions are achieved by keeping the spatial feature map from the second last layer in the video network instead of the 1-D feature vector after the last pooling layer as in the above settings. In this way, the merged features would account for not only global information, but also local structural patterns (anatomies). Here, the audio-video fusion is achieved by element-wise dot product between the audio feature vector and each of the visual feature vector in the visual feature map, followed by global average pooling and fully-connected layers in an attention-like aggregation manner into a single vector before being fed into the loss functions. The element-wise dot product is used to measure the similarity between the audio feature and the visual features, aiming to activate the high-response regions on the spatial visual feature map. As illustrated in Figure 2d, by using the spatial-aware fusion strategy, the multi-modal features are able to capture local spatial information instead of global semantics only (as in Figure 2d-I).

In this case, the final fused feature is a three-dimensional tensor (instead of one-dimensional vector) with spatial information that represents the spatial relationship between the visual and audio signals. The corresponding loss function used for training is then defined as:

$$L_{base}(V_i^s, A_i) + L_{contra}(V_i^s, A_i), \quad (3)$$

where V_i^s is the 3-D visual feature with spatial information.

Adding textual information The speech audio data modality provides auxiliary supportive information to assist in video understanding and representation learning, but according to our observations, in our application, it can also bring ambiguity in correlation modelling. Specifically, the audio signal is not generated from the visual source and as a result, may have a weak (or even no) correlation with video content. For example, the content of the audio is often a conversation between the sonographer and the subject while the video being shown on the screen is that of fetal US content. There are even different kinds of background sounds that can exist in the audio data (we have observed ambulance sirens, people coming into the




Figure 3. Audio data cleaning. (a) Pipeline for the proposed audio data cleaning. The input audio signal (the left red waveform) results in an enhanced audio signal (the bottom-right green waveform) from the sonographer accompanied by a corresponding text segment (the bottom-left text). (b) Example audio data samples *before* (top) and *after* (bottom) the proposed audio data pre-processing.

scanning room, etc). These challenges prevent the model from learning accurate and robust multi-modal representations. To this end, we propose to leverage additional information along with the audio signal, by including a third data modality – text. We automatically labelled the speech data using automatic speech recognition (ASR). Textual transcriptions are acquired by applying an ASR model to the speech audio data. In this way, we can use the textual feature embedding to help the multi-modal US representation learning, and this also allows us to focus on key information (*e.g.* anatomy-related terminology) while eliminating unrelated information by keyword spotting. The basic idea is illustrated in Figure 2c. Specifically, we use the word2vec model [16] followed by two fully-connected layers to project textual information into word embeddings.

However, due to the poor performance of the ASR model when applying to unseen US data, the ASR transcribed text cannot be directly leveraged as input for the model training. Apart from the existing challenges in audio, the ASR transcription has many errors (*e.g.* non-meaningful English), which may further confuse the model during the joint representation learning. Therefore, we propose approaches to address these issues (*i.e.* noise in audio, ambiguity in transcribed text, *etc.*), as will be elaborated on in the following sections.

Selective Information Gating

As mentioned above, the transcribed raw textual data is inappropriate to be directly fed into the text embedding model. In this section, we introduce a new selective information gating (SIG) module to filter out the undesired information and only preserve the key information that we are interested in in a selective manner. The proposed SIG module is designed in two variants: outlier filtering and keywords spotting, as illustrated in Figure 2e and detailed below.

Filtering outliers The first SIG approach is at a relatively coarse level, where only outliers are removed. Here we consider those wrongly transcribed words *e.g.* invalid English words and blank text segments (zero padding) as outliers. Whenever such outliers are detected in the input text, the corresponding embeddings are removed and thus do not contribute to the final correlation modelling. The wrongly transcribed words are effectively identified before being fed into the SIG module, and the module simply filters them out to remove their possible influence on model training. These words have been identified through the tokenizer used with the word2vec embedding layer, and such ‘unknown words’ that do not exist in the word2vec dictionary would be represented with a $\langle \text{unk} \rangle$ token. Then it is the word2vec embedding of this token that gets filtered out by the SIG module. More specifically, the SIG module consists of a convolutional layer with 1×1 kernel size followed by a fully-connected layer, which outputs a weighting vector to combine with the word embeddings, as shown in Figure 2e-I. The projection layers are fully-connected layers to project the word2vec embeddings to combine with the SIG module output. The final filtered text embedding is used to optimise the multi-modal learning objective.

Keywords spotting Even when the ASR model is able to correctly recognise all the remaining speech from the SIG output, there can be some content that is not related to the video. For instance, another medical professional enters the room and has a conversation with the sonographer that is unrelated to the US scan. In this case, modelling the text (as well as the audio) with the video would confuse the model. To reduce this potential source of error, we introduce a *keywords dictionary* to the SIG module, so that the SIG module is aware of the keywords and select the corresponding information (Figure 2e-II). Specifically, the SIG module only focuses on the keywords within the dictionary when predicting the weights for the gating operation, in a compare and matching manner by filtering out and only selecting those words exist in the dictionary while depressing the features of other words. The keywords in the dictionary are constructed empirically with domain knowledge from the clinicians.

Multi-Modal Optimisation

To optimise the multi-modal joint representation learning, we propose to use a combined loss function that consists of correspondence modelling and contrastive learning for the video-audio-text modalities. More specifically, according to the definitions in equation (1) and equation (2), the final joint loss function for multi-modal optimisation is defined as

$$\mathcal{L} = \alpha \cdot L_{base}(V_i, A_i) + \beta \cdot L_{base}(V_i, T_i) + \gamma \cdot L_{contra}(V_i, A_i) + \delta \cdot L_{contra}(V_i, T_i), \quad (4)$$

where T_i is the text sample and $\{\alpha, \beta, \gamma, \delta \mid \alpha + \beta + \gamma + \delta = 1\}$ are weighting parameters to balance the contribution of each loss term.

Results

In this section, we first describe the detailed experimental settings including the data and evaluation strategies. We will also compare with other alternative or state-of-the-art solutions, with isolated ablation studies to validate the effectiveness of the proposed approach.

Data

Full-length routine fetal second-trimester US videos were acquired by expert sonographers as part of the PULSE study. As there is no public dataset available with the same clinical setting as ours to our knowledge, here the experiments are performed on our PULSE data. Sonographer audio recordings that include audible vocal descriptions of the performed scans were collected with the videos, creating an audio-visual dataset. In total, there are 81 scans with such audio and video data and each scan contains around 55,000 frames. Each whole scan was divided into smaller video clips to feed into the networks, resulting in a total of 73,681 clips. The speech audio data was transcribed via ASR, allowing for the inclusion of the third modality, text. As mentioned above, the text was further processed by filtering unrelated words and focusing on the keywords for the text embedding extraction. The keywords dictionary consists of anatomical terminologies within the fetal US scenario. When sampling the positive and negative training pair, for each video sample at a particular timestamp t , the corresponding audio and text are sampled in a neighbour region with the time range of 0.6 s (*i.e.* $t \pm 0.3s$ of the video sample at t). Before feeding into the audio network, the original audio data was converted to the 2-dimensional log-spectrogram format with size 256×256 using a short-time Fourier transform with frequency bands of 256, window length 10ms and hop length of 5ms. For the aforementioned downstream tasks, we used an additional set of 135 scans and conducted three-fold cross-validation (by randomly selecting 90 scans for training in each fold).

Audio Data Cleaning

An audio cleaning step was implemented as illustrated in Figure 3a. The raw audio data contains unwanted background signals (including the sound of US machine keyboard actions and the room air conditioner fan among other background sounds) as well as irrelevant, undesired speech. The first step to reduce undesired audio is to apply audio denoising followed by voice activity detection, in which the former is achieved by speech enhancement [7] and the latter by applying VADNet [20]. That makes it possible to differentiate speech segments from noise segments as well as seconds of silence.

After background noise reduction, we further separate the audio signal into different speakers using speaker diarization. Speaker diarization is a speech processing technique that aims to separate an audio recording into two (or more) speech segments where each segment contains a single unique speaker. Specifically, we employ the pyannote-audio [5] framework to perform speaker diarization and assume that the speaker with the most speech segments is the sonographer. By only preserving the sonographer’s speech and disposing of the speech of other people, we have a “clean” speech dataset suitable for analysis. Examples of the audio data before the pre-processing and after the pre-processing are presented in Figure 3b for reference.

Finally, transcribing the sonographer speech segments with a fine-tuned wav2vec-based automatic speech recognition (ASR) model [4], we identify the speech segments where anatomically relevant information is being spoken by comparing the transcribed text with an existing dictionary of anatomically relevant keywords.

Assessment and Evaluation Criteria

After learning the multi-modal representations, we need an assessment tool to evaluate the quality of the learned representations and their applications. Following prior works [2, 12], we use several downstream tasks to assess the quality of these representations, in a transfer learning manner. Here, in our fetal US scan setting, three downstream tasks are chosen: standard plane detection, eye-gaze saliency prediction, and audio-guided visual localisation. Through these three tasks, we not only evaluate the learned visual representations (as in prior works) but also evaluate the quality of the joint representation (by the third task). The standard plane detection is a classification problem, where given an US image, the task aims to classify it into an anatomical category. The eye-gaze saliency prediction is a regression task that tries to predict the eye-gaze attention

regions that the sonographer is focusing on. The visual localisation task aims to localise the anatomy regions in the US video according to the indication from the corresponding audio data. The first two downstream tasks focus more on the visual representations, while the third one needs the model to have a thorough multi-modal understanding so that it can successfully localise the audio-guided visual anatomies. For quantitative assessment, we leverage the commonly used *precision*, *recall* and *F1* score to evaluate the standard plane detection task; the Kullback-Leibler divergence (*KL*), normalised scanpath saliency (*NSS*), area under ROC curve (*AUC*), Pearson’s correlation coefficient (*CC*) and similarity (*SIM*) [6] for the saliency prediction task; as there is no localisation annotations available in our dataset and it is hard to acquire accurate pixel-level annotations, it is infeasible to present a quantitative evaluation on the third task. Therefore, we consider eye-gaze saliency to be a close enough approximation to expert-made annotations as the eye gaze data comes from the experts themselves. Eye-gaze saliency has been used in other work [1, 19] as well in an effort to tie language with visual regions of interest. We provide a corresponding quantitative evaluation in this case.

We also found that the saliency prediction in the clinical setting is different from that in natural vision scenarios [6], as we focus on more fine-grained anatomical local structures in medical imaging. As a result, in addition to the aforementioned metrics, we further design a new metric tailored for eye-gaze saliency assessment in the clinical setting. This metric is designed to model the compactness of the predicted saliency region, as well as the intensity of the highly-confident response. Specifically, assuming the area of the whole saliency regions to be $\alpha_S = \|I_S\|_0$ (I_S is the normalised saliency map) and the highly-confident response region as $\alpha_H = \|I_{HS}\|_0$ where $I_{HS} = I_S > thres$, the high response intensity is defined as $\eta_H = \alpha_H / \alpha_S$. Additionally, we also consider the correctness of the predicted saliency when measuring its compactness, by using the intersection-over-union (IoU) metric (IoU=intersection(I_S, \hat{I}_S)/union(I_S, \hat{I}_S)). The proposed compactness metric is then defined as

$$Comp = \text{IoU} \cdot \frac{\eta_H}{\alpha_S}. \quad (5)$$

Implementation Details

We follow our preliminary work [11] for the baseline network design and the parameter settings. The ResNeXt50 [21] was used as the backbone with squeeze-and-excitation module and dilated convolutions for the video branch. To take images/frames as input for the downstream tasks, following [11], within an interval of a video clip, we randomly sample a few frames (2 in this work due to the memory limitation) to represent the video clip, and each of them is fed into the encoder for the feature extraction. The corresponding features are concatenated subsequently as the video feature. Following [11], the audio branch shares the same network architecture as the video branch but is optimised separately and with different inputs. The speech audio was extracted with a 0.6s interval and resampled to 24kHz before being converted to a 2D log-spectrogram representation (of size 256×256) using a short-time Fourier transform. Then this log-spectrogram is fed into the audio net for the following representation learning. The text branch and audio branch are designed in a similar manner except that we used the word2vec model [16] for text embedding extraction, while the remaining joint modelling with the video embedding is performed in the same way. We also tried a more advanced text embedding model (BERT [8]) in replacement of the word2vec model, but did not observe any noticeable improvement. We suspect the reason lies in the relatively simpler context of our text data, while the BERT model was designed to address the ambiguous context situation where the same word can have more than one meaning depending on the context. Most of our words are high-level anatomical (e.g. ‘bones’, ‘heart’, ‘arms’, ‘legs’, ‘blood’, ‘kidneys’, and ‘elbow’), and are adequately covered by a general corpus. BioBERT is an alternative to BERT that has been trained on biomedical corpora. BioBERT has been shown to work well for specialist biomedical terms such as ‘antimicrobial’ and ‘transcriptional’ [15]. However, despite the medical context in which our data is acquired, our vocabulary is more conversational because it is spoken between a clinical professional and the patient. It did not seem necessary to use such a specialised embedding model. In terms of optimisation, we use a similar combined objective as in the aforementioned audio-video modelling, *i.e.* correspondence modelling with contrastive learning. The weighting parameters in the final loss (equation (4)) are empirically set to be equal after normalising to the same scale. We also did an analysis on the weighting parameters by setting different values to each of the parameters, and found that the model performs the best when they are set to be equal. The whole model was trained end-to-end with the stochastic gradient descent (SGD) optimiser and the learning rate was initialised as 10^{-3} and decayed by the scale of 10 for every 20 epochs. The batch size was set to 40. The visual signal was scaled to 256×256 and centre cropped to 224×224 to feed into the GPU. We implemented with the PyTorch deep learning framework on a workstation equipped with NVIDIA GPU cards (Titan V and Tesla V100).

Representation Quality Assessment

Standard plane detection As mentioned above, the quality of the learned multi-modal representations is assessed by transfer learning to downstream tasks. First we perform evaluation on the standard plane detection (SPD) task. The 14 categories used in this task are: three-vessel tracheal view of the heart (3VT), four-chamber view of the heart (4CH), left ventricular outflow tract of the heart (LVOT), right ventricular outflow tract of the heart (RVOT), transventricular plane of the brain (BrainTv.),

Table 1. Evaluation results on standard plane detection (mean±std.[%]). Best performance is marked in **bold**. Note the methods on the right side are fully-supervised using external manual annotations. All the three metrics are the higher the better.

	Rand.Init.	Video	V-A Base	V-T Base	Ours	ImageNet Init.	US-Sup
Precision	70.4±2.3	71.9±2.0	72.7±1.8	75.6±1.4	77.2±1.4	74.6±1.8	87.1±1.9
Recall	64.9±1.6	71.7±3.5	73.3±2.4	75.4±1.6	77.3±1.8	71.2±1.9	81.0±2.7
F1-score	67.0±1.3	71.5±2.4	72.6±1.7	75.1±0.2	76.8±1.0	72.5±1.8	82.7±1.6

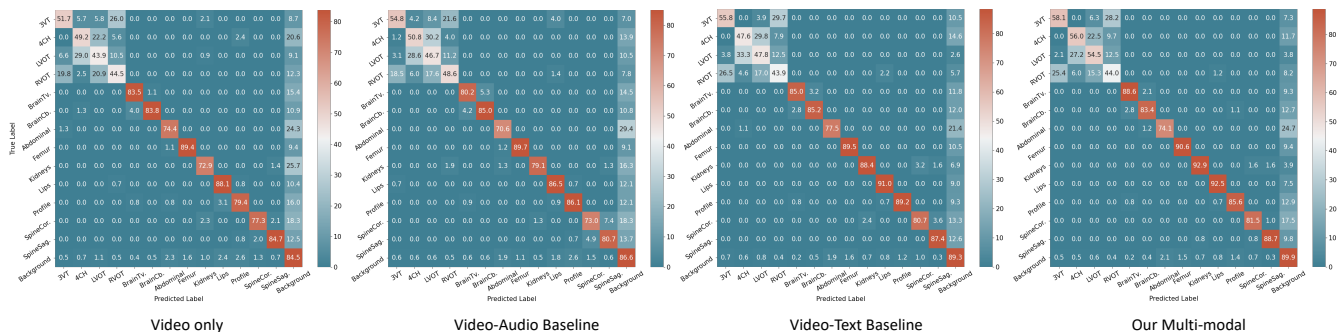


Figure 4. Confusion matrices of the representation learning methods on the SPD task.

transcerebellar plane of the brain (BrainCb.), abdomen, femur, kidneys, lips, profile, spine in the coronal plane (SpineCor.), spine in the sagittal plane (SpineSag.) and background. The performance of this classification task is reported in Table 1, where we include the results of using random initialisation, only video modality, video-audio baseline [11], video-text baseline and our multi-modal approach (*Ours*) as initialisation for the SPD model. Here the *Video* refers to the self-supervised representation learning approach (with a frame order prediction pretext task) using only video data for pre-training; and the *V-A Base* refers to the video-audio self-supervised representation learning approach proposed in [11]. The *V-T Base* is the baseline using only the video and text branch (without the proposed keywords spotting SIG module) as shown in Figure 2c. Additionally, we also report the results of two fully-supervised models which were trained with manual annotated ground-truth labels: *ImageNet Init.* for using the ImageNet [18] dataset as training data, and *US-Sup* for using US data with expert annotations as supervision for the same standard plane detection task during pre-training. From the quantitative results shown in Table 1, we can see that by using the additional audio data, the performance of SPD can be improved compared to using video only. When using the proposed multi-modal approach, the performance is further improved by a large margin (*e.g.* 72.7→77.2 for the precision). It can be seen that when only using the video and text branches, the model also performs quite well, suggesting the effectiveness of the data processing and text-related modules as well as the joint modelling. It is interesting to see that even though the *ImageNet Initialisation* was trained with human annotations in a fully-supervised manner, the proposed approach still performs much better. As *US-Sup* was pre-trained with the exact same task of SPD and with human annotations, it can be considered as an upper bound. To better understand how the representation learning methods perform in more detail, we also present the confusion matrix over all the 14 categories, as shown in Figure 4. From the confusion matrix, it can be seen that the cardiac-related categories (*i.e.* 3VT, 4CH, LVOT, RVOT) are relatively poorly performing; it is well known that these are hard even for a human expert. However, when we compare the *Video-Audio Baseline* with the *Video only*, the performance of these categories is greatly improved, suggesting the effectiveness of the auxiliary audio modality. Building upon that, the proposed multi-modal spatial aware approach further improved almost all the categories significantly. For instance, *LVOT*: 46.7→54.5 and *Kidneys*: 79.1→92.9. This fine-grained analysis further validates the effectiveness of the proposed method.

Eye-gaze saliency prediction As in our PULSE dataset, the eye-gaze tracking information of the sonographers was captured while performing US, in addition to the above SPD task; this allows us to take the eye-gaze saliency prediction as an additional downstream task to evaluate the quality of the learned representations. Similar to the SPD task, the pre-trained model was loaded to the saliency prediction model followed by fine-tuning. The only difference is that the last layers were modified so that the model is able to predict a spatial saliency map instead of a single category. The quantitative performance is reported in Figure 5, with the evaluation metrics mentioned above. Again, we can see that with the addition of the audio modality, the model performs much better when compared to using only video data. The textual information also helps the representation learning, with similar performance as the video-audio baseline and outperforming the single-modal approach. When jointly trained with video, audio and text, the proposed multi-modal approach further boosts the performance by a large margin.




Figure 5. Evaluation on eye-gaze saliency prediction. Quantitative performance (a) evaluated by conventional metrics, and (b) with the newly proposed *Comp* metric. ↓ indicates the lower the better while ↑ means the higher the better. (c) Qualitative performance on eye-gaze saliency prediction.

Different from the results for the standard plane detection task (Table 1), in this regression task we can see that our method outperforms the supervised settings (*ImageNet Init.* and *US-Sup*), suggesting that our approach has better generalisability. We also present qualitative examples in Figure 5c, in which the predicted saliency map is overlaid to the original US image. This qualitative comparison further validates the effectiveness of the proposed method. An interesting observation is that, when comparing our approach with the video-audio baseline, we can see (in the qualitative results) that our predicted saliency is more compact and closer to the human performance, though for some of the quantitative metrics shown in Figure 5a, the difference is not that significant (*Video-audio baseline* vs. *Ours*). As clarified above, this is mainly due to the fact that the metrics in Figure 5a are proposed for natural images instead of medical imaging. When we conduct a further evaluation by using the newly proposed *Comp* metric (as in Figure 5b), we can see that the proposed approach has a much better compactness performance, aligning well with the qualitative results.

Ablation Study

In order to analyse the effectiveness of each component in the proposed framework, we conducted a controlled experiment (*i.e.* ablation study), as shown in Figure 6. Performance on the above two downstream tasks is reported with the corresponding evaluation metrics. Here we chose the baseline as the very basic audio-video correspondence modelling (Figure 2a). Then we compared this to the performance of adding cross-modal contrastive learning (*CM. Contra.*), the aforementioned video-audio baseline, the spatial awareness, the additional textual information, only using visual and text path for training (*w/o Audio*) and the SIG module. We can see that each component contributes to the performance gain of our proposed approach, validating their effectiveness.

Audio-Guided Visual Anatomy Localisation

We have shown that the proposed multi-modal audio-visual modelling framework is able to perform audio-guided visual anatomy localisation with no extra human interventions. In this section, we showcase how this can be achieved. During the training of the multi-modal representation learning framework, audio (as well as derived text) embedding perceive the

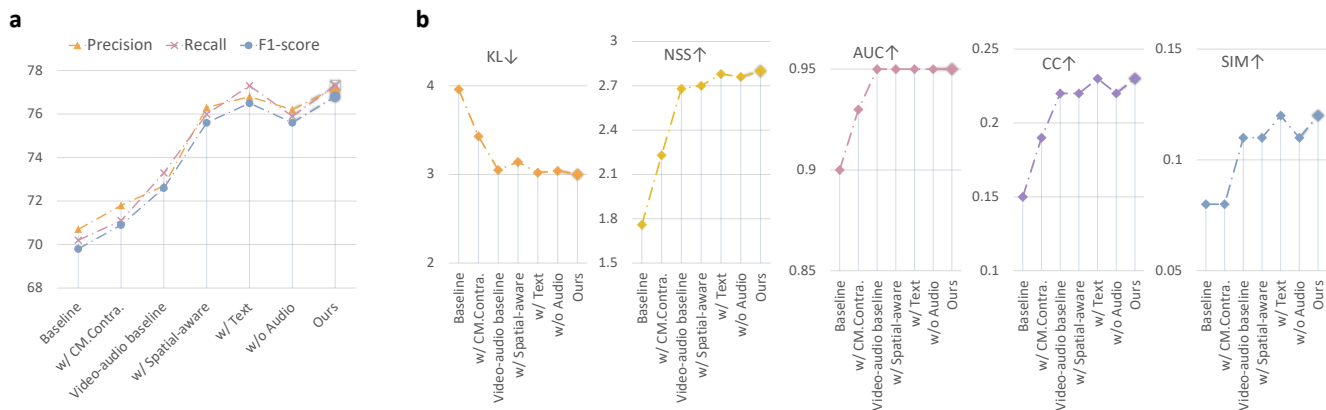


Figure 6. Ablation study on each of the proposed strategies for two downstream tasks. (a) Standard plane detection. (b) Eye-gaze saliency prediction, where ↓ indicates the lower the better while ↑ means the higher the better.

spatial information of the visual data well, by the fusion module. As a result, after the training, given input audio the model is able to highlight the high-response visual regions and such regions are usually related to anatomies due to the training target. Note that instead of visualising landmarks, our model shows the association between the visual signal and the audio input, trying to highlight the corresponding anatomical structures. In Figure 7a, we showcase examples of this audio-guided visual anatomy localisation task. We can see that our model is able to correctly localise the anatomies that the speech audio contains, *e.g. heart, kidneys etc.*. This demonstrates the quality of the learned joint multi-modal representations and also illustrates a useful application towards anatomy localisation. The strategy for the audio-visual anatomy localisation task is the same setting as in our work mentioned above, *i.e.* the speech audio data is fed into the pretrained audio net together with the video data captured from the ultrasound machine being fed into the video net. Then the model is able to highlight the high-response visual anatomical regions accordingly. Under our experimental settings, the time to process a 244×244 frame and generate the corresponding response map is 0.08s, *i.e.* over 12 FPS (frames per second), which is acceptable for applications. As mentioned above, it is infeasible to acquire accurate clinical annotations for this task and we instead seek help from the eye-gaze saliency data as a reference to provide an approximate quantitative evaluation here. Specifically, we compare the localisation map of the prediction from our model and the eye-gaze saliency from the captured gaze data, and measure their similarity on a separate unseen test set with both audio and gaze data. The results are shown in Figure 7c, using the evaluation metrics of KL, CC, SIM and Comp (NSS and AUC are discarded as in this task we are not measuring human fixations). For comparison, we also include a saliency prediction method that was trained in a fully-supervised manner with ground-truth supervision from eye-gaze saliency data, and applied to the same test data. Note our method in this audio-guided visual localisation task never saw any gaze/saliency data but was only guided by the audio signal. From the results we can see that even without seeing any eye gaze data, our method still performs well and better (3 out of 4 metrics) than the solution that was trained with eye-gaze data as ground-truth supervision. As described above, the content of the speech audio is not always related to anatomical regions. In the case where the audio does not relate to any anatomy, the model may fail to give meaningful localisation. Examples of such “failure cases” are shown in Figure 7b for reference. Interestingly, though the audio content does not relate to any anatomy, the model still tries to reason about some localisation, *e.g.* the model is looking at the anatomy and its neighbours to try to ‘figure out’ the “six weeks” narration.

Discussion

Limitations Although this work made the first attempt towards learning multi-modal representations in real-world clinical settings and demonstrated meaningful representations, there are some limitations that may prevent the proposed approach from being applied more generally. In this work, we only validated the proposed model on clinical fetal US data. The proposed framework is general and there are no assumptions which make it obstetrics specific, however it has not been validated in other clinical settings as yet to prove generalisability. Another limitation is that the current model is designed to localise *anatomy-related* regions and not more general concepts (as shown in Figure 7b). The main reason is that the most consistent audio-video pairs are related to anatomy in our case; if other well-paired training data for other concepts were available, this limitation could be addressed accordingly.

Potential impact and benefits The study was conducted using data from one hospital site in Oxford, thus the corresponding conclusions may not be identical for data from other sites. However, data was collected in a real-world setting; this also explains

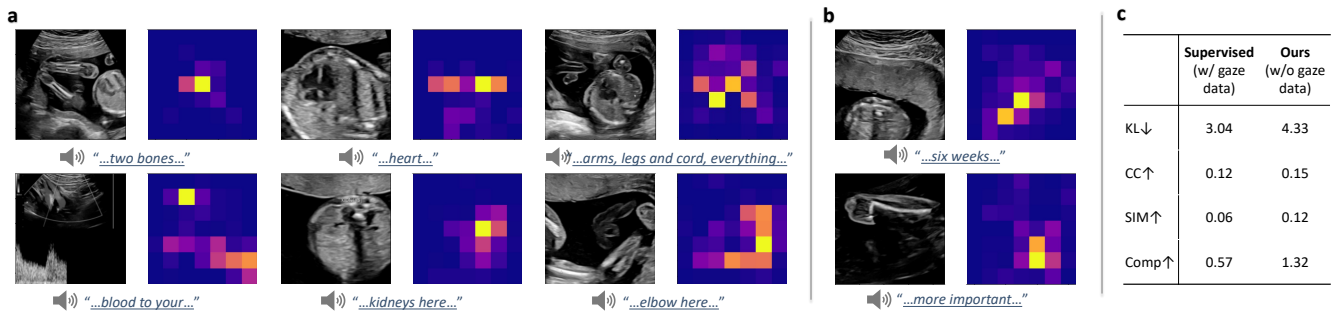


Figure 7. Performance on audio-guided visual anatomy localisation. (a) Performance on normal cases where audio is related to the visual anatomy, while in (b) the audio is unrelated to anatomy. In each sample, the left shows the input US with the corresponding speech audio (transcription) below; the right side shows the estimated localisation map from the proposed method. Higher response is in yellow colour while lower in blue colour. (c) Quantitative evaluation of the audio-guided visual localisation task.

the presence of audio signals that are not relevant to anatomical regions, for example conversations between the sonographer and prospective parents. Thus, we believe that the model should apply to most similar settings, though including more data from other clinical environments could strengthen the model. There are several strengths of this work. We have made the first attempt towards multi-modal clinical representation learning with a particular focus on video-audio data modalities. This enables us to have a better and more thorough understanding of the clinical data and eliminates the requirement of human expert annotations. In some clinical settings, including the UK where this study was conducted, sonographers will communicate their findings and provisional information to the patient, thereby providing an additional modality to the data science of sonography. And sonographers usually impart such information during the scan, and confirming the normality of views or absence of abnormalities. The proposed model can also be utilised to help localise anatomical regions of interest, by only using speech audio, *i.e. talk to the screen, and show from tell*. This has both scientific and clinical contributions. Data annotation can be made much easier by only verbally describing the regions that need to be labelled. The proposed model may be useful to help the training of new sonographers.

Conclusion

In this work, we presented a study on multi-modal representation learning in clinical settings, with a particular focus on audio-visual data. A new multi-modal framework has been proposed with additional textual data and a selective gating module. Extensive experimental analysis on a large-scale clinical dataset shows that the proposed approach can learn meaningful representations without additional human expert intervention. The learned representations were shown to be effective in several downstream clinical tasks. We also introduced a new clinical application, audio-guided visual localisation, which may be useful to both clinicians and patients in video interpretation.

Acknowledgements

The authors would like to thank the support from the EPSRC Programme Grant Visual AI EP/T028572/1, the ERC Project PULSE ERC-ADG-2015 694581, and the support of NVIDIA Corporation with the donation of the GPU. A.T.P. is supported by the Oxford Partnership Comprehensive Biomedical Research Centre with funding from the NIHR Biomedical Research Centre (BRC) funding scheme.

References

1. Mohammad Alsharid, Yifan Cai, Harshita Sharma, Lior Drukker, Aris T Papageorghiou, and J Alison Noble. Gaze-assisted automatic captioning of fetal ultrasound videos using three-way multi-modal deep neural networks. *Medical Image Analysis*, 82:102630, 2022.
2. Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–617, 2017.
3. Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European conference on computer vision (ECCV)*, pages 435–451, 2018.
4. Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *arXiv preprint arXiv:2006.11477*, 2020.

5. Hervé Bredin, Ruiqing Yin, Juan Manuel Coria, Gregory Gelly, Pavel Korshunov, Marvin Lavechin, Diego Fustes, Hadrien Titeux, Wassim Bouaziz, and Marie-Philippe Gill. Pyannote. audio: neural building blocks for speaker diarization. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7124–7128. IEEE, 2020.
6. Zoya Bylinskii, Tilke Judd, Aude Oliva, Antonio Torralba, and Frédo Durand. What do different evaluation metrics tell us about saliency models? *IEEE transactions on pattern analysis and machine intelligence*, 41(3):740–757, 2018.
7. Alexandre Defossez, Gabriel Synnaeve, and Yossi Adi. Real time speech enhancement in the waveform domain. In *Interspeech*, 2020.
8. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
9. Lior Drukker, Harshita Sharma, Richard Droste, Mohammad Alsharid, Pierre Chatelain, J Alison Noble, and Aris T Papageorghiou. Transforming obstetric ultrasound into data science using eye tracking, voice recording, transducer motion and ultrasound video. *Scientific Reports*, 11(1):1–12, 2021.
10. Zeyu Fu, Jianbo Jiao, Robail Yasrab, Lior Drukker, Aris T. Papageorghiou, and Alison Noble. Anatomy-aware contrastive representation learning for fetal ultrasound. In *European Conference on Computer Vision Workshop*, 2022.
11. Jianbo Jiao, Yifan Cai, Mohammad Alsharid, Lior Drukker, Aris T Papageorghiou, and J Alison Noble. Self-supervised contrastive video-speech representation learning for ultrasound. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 534–543. Springer, 2020.
12. Jianbo Jiao, Richard Droste, Lior Drukker, Aris T Papageorghiou, and J Alison Noble. Self-supervised representation learning for ultrasound video. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 1847–1850. IEEE, 2020.
13. Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
14. Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
15. Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020.
16. Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
17. Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12475–12486, 2021.
18. Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
19. Yusuke Sugano and Andreas Bulling. Seeing with humans: Gaze-assisted neural image captioning. *arXiv preprint arXiv:1608.05203*, 2016.
20. Johannes Wagner, Dominik Schiller, Andreas Seiderer, and Elisabeth André. Deep learning in paralinguistic recognition tasks: Are hand-crafted features still relevant? In *Interspeech*, 2018.
21. Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.