
CLEAR: A Dataset for Compositional Language and Elementary Acoustic Reasoning

Jerome Abdelnour
NECOTIS, ECE Dept.
Sherbrooke University
Québec, Canada
Jerome.Abdelnour
@usherbrooke.ca

Giampiero Salvi
KTH Royal Institute
of Technology
EECS School
Stockholm, Sweden
giampi@kth.se

Jean Rouat
NECOTIS, ECE Dept.
Sherbrooke University
Québec, Canada
Jean.Rouat
@usherbrooke.ca

Abstract

We introduce the task of acoustic question answering (AQA) in the area of acoustic reasoning. In this task an agent learns to answer questions on the basis of acoustic context. In order to promote research in this area, we propose a data generation paradigm adapted from CLEVR [11]. We generate acoustic scenes by leveraging a bank elementary sounds. We also provide a number of functional programs that can be used to compose questions and answers that exploit the relationships between the attributes of the elementary sounds in each scene. We provide AQA datasets of various sizes as well as the data generation code. As a preliminary experiment to validate our data, we report the accuracy of current state of the art visual question answering models when they are applied to the AQA task without modifications. Although there is a plethora of question answering tasks based on text, image or video data, to our knowledge, we are the first to propose answering questions directly on audio streams. We hope this contribution will facilitate the development of research in the area.

1 Introduction and Related Work

Question answering (QA) problems have attracted increasing interest in the machine learning and artificial intelligence communities. These tasks usually involve interpreting and answering text based questions in the view of some contextual information, often expressed in a different modality. Text-based QA, use text corpora as context ([19, 20, 17, 9, 10, 16]); in visual question answering (VQA), instead, the questions are related to a scene depicted in still images (e.g. [11, 2, 25, 7, 1, 23, 8, 10, 16]). Finally, video question answering attempts to use both the visual and acoustic information in video material as context (e.g. [5, 6, 22, 13, 14, 21]). In the last case, however, the acoustic information is usually expressed in text form, either with manual transcriptions (e.g. subtitles) or by automatic speech recognition, and is limited to linguistic information [24].

The task presented in this paper differs from the above by answering questions directly on audio streams. We argue that the audio modality contains important information that has not been exploited in the question answering domain. This information may allow QA systems to answer relevant questions more accurately, or even to answer questions that are not approachable from the visual domain alone. Examples of potential applications are the detection of anomalies in machinery where the moving parts are hidden, the detection of threatening or hazardous events, industrial and social robotics.

Current question answering methods require large amounts of annotated data. In the visual domain, several strategies have been proposed to make this kind of data available to the community [11, 2, 25, 7]. Agrawal et al. [1] noted that the way the questions are created has a huge impact on what information a neural network uses to answer them (this is a well known problem that can arise with

| Question type | Example | Possible Answers | # |
|---------------|---|---|----|
| Yes/No | Is there an equal number of <i>loud cello</i> sounds and <i>quiet clarinet</i> sounds? | yes, no | 2 |
| Note | What is the note played by the <i>flute</i> that is <i>after the loud bright D</i> note? | A, A#, B, C, C#, D, D#, E, F, F#, G, G# | 12 |
| Instrument | What instrument plays a <i>dark quiet</i> sound in the <i>end</i> of the scene? | cello, clarinet, flute, trumpet, violin | 5 |
| Brightness | What is the brightness of the <i>first clarinet</i> sound? | bright, dark | 2 |
| Loudness | What is the loudness of the <i>violin</i> playing after the <i>third trumpet</i> ? | quiet, loud | 2 |
| Counting | How many other sounds have the same brightness as the <i>third violin</i> ? | 0–10 | 11 |
| Absolute Pos. | What is the position of the <i>A#</i> note playing after the <i>bright B</i> note? | } first–tenth | 10 |
| Relative Pos. | Among the <i>trumpet</i> sounds which one is a <i>F</i> ? | | |
| Global Pos. | In what part of the scene is the <i>clarinet</i> playing a <i>G</i> note that is <i>before</i> the <i>third violin</i> sound? | beginning, middle, end (of the scene) | 3 |
| Total | | | 47 |

Table 1: Types of questions with examples and possible answers. The variable parts of each question is emphasized in bold italics. In the dataset many variants of questions are included for each question type, depending on the kind of relations the question implies. The number of possible answers is also reported in the last column. Each possible answer is modelled by one output node in the neural network. Note that for absolute and relative positions, the same nodes are used with different meanings: in the first case we enumerate all sounds, in the second case, only the sounds played by a specific instrument.

all neural network based systems). This motivated research [23, 8, 11] on how to reduce the bias in VQA datasets. The complexity around gathering good labeled data forced some authors [23, 8] to constrain their work to yes/no questions. Johnson et al. [11] made their way around this constraint by using synthetic data. To generate the questions, they first generate a semantic representation that describes the reasoning steps needed in order to answer the question. This gives them full control over the labelling process and a better understanding of the semantic meaning of the questions. They leverage this ability to reduce the bias in the synthesized data. For example, they ensure that none of the generated questions contains hints about the answer.

Inspired by the work on CLEVR [11], we propose an acoustical question answering (AQA) task by defining a synthetic dataset that comprises audio scenes composed by sequences of elementary sounds and questions relating properties of the sounds in each scene. We provide the adapted software for AQA data generation as well as a version of the dataset based on musical instrument sounds. We also report preliminary experiments using the FiLM architecture derived from the VQA domain.

2 Dataset

This section presents the dataset and the generation process¹. In this first version (version 1.0) we created multiple instances of the dataset with 1000, 10000 and 50000 *acoustic scenes* for which we generated 20 to 40 questions and answers per scene. In total, we generated six instances of the dataset. To represent questions, we use the same semantic representation through functional programs that is proposed in [11, 12].

2.1 Scenes and Elementary Sounds

An acoustic scene is composed by a sequence of *elementary sounds*, that we will call just sounds in the following. The sounds are real recordings of musical notes from the Good-Sounds database [3]. We use five families of musical instruments: cello, clarinet, flute, trumpet and violin. Each recording of an instrument has a different musical note (pitch) on the MIDI scale. The data generation process, however, is independent of the specific sounds, so that future versions of the data may include speech, animal vocalizations and environmental sounds. Each sound is described by an n-tuple [Instrument family, Brightness, Loudness, Musical note, Absolute Position, Relative Position, Global Position, Duration] (see Table 1 for a summary of attributes and values). Where Brightness can be either bright or dark; Loudness can be quiet or loud; Musical note can take any of the 12 values on the fourth octave of the Western chromatic scale². The Absolute Position gives the position of the sound within the acoustic scene (between first and tenth), the Relative Position gives the position of a sound relatively to the other sounds that are in the same category (e.g. “the third cello sound”). Global Position refers

¹Available at <https://github.com/IGLU-CHISTERA/CLEAR-dataset-generation>

²For this first version of CLEAR the cello only includes 8 notes: C, C#, D, D#, E, F, F#, G.

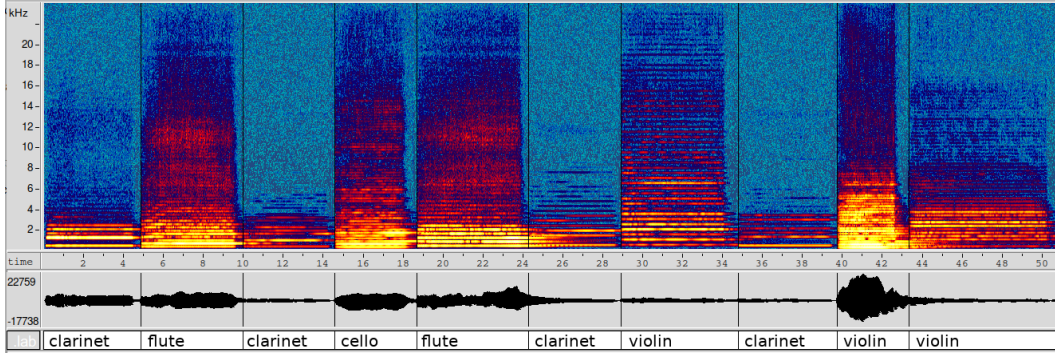


Figure 1: Example of an acoustic scene. We show the spectrogram, the waveform and the annotation of the instrument for each elementary sounds. A possible question on this scene could be "What is the position of the flute that plays after the second clarinet?", and the corresponding answer would be "Fifth". Note that the agent must answer based on the spectrogram (or waveform) alone.

to the approximate position of the sound within the scene and can be either beginning, middle or end.

We start by generating a *clean* acoustic scene as following: first the encoding of the original sounds (sampled at 48kHz) is converted from 24 to 16 bits. Then silence is detected and removed when the energy, computed as $10 \log_{10} \sum_i x_i^2$ over windows of 100 msec, falls below -50 dB, where x_i are the sound samples normalized between ± 1 . Then we measure the perceptual loudness of the sounds in dB LUFS using the method described in the ITU-R BS.1770-4 international normalization standard [4] and implemented in [18]. We attenuate sounds that are in an intermediate range of -24 dB LUFS and -30.5 dB LUFS by -10 dB, to increase the separation between loud and quiet sounds. We obtain a bank of 56 elementary sounds. Each clean acoustic scene is generated by concatenating 10 sounds chosen randomly from this bank.

Once a clean acoustic scene has been created it is post-processed to generate a more difficult and realistic scene. A white uncorrelated uniform noise is first added to the scene. The amplitude range of the noise is first set to the maximum values allowed by the encoding. Then the amplitude is attenuated by a factor f randomly sampled from a uniform distribution between -80 dB and -90 dB ($20 \log_{10} f$). The noise is then added to the scene. Although the noise is weak and almost imperceptible to the human ear, it guaranties that there is no pure silence between each elementary sounds. The scene obtained this way is finally filtered to simulate room reverberation using SoX³. For each scene, a different room reverberation time is chosen from a uniform distribution between [50ms, 400ms].

2.2 Questions

Questions are structured in a logical tree introduced in CLEVR [11] as a *functional program*. A functional program, defines the reasoning steps required to answer a question given a scene definition. We adapted the original work of Johnson *et al.* [11] to our acoustical context by updating the function catalog and the relationships between the objects of the scene. For example we added the *before* and *after* temporal relationships.

In natural language, there is more than one way to ask a question that has the same meaning. For example, the question "Is the cello as loud as the flute?" is equivalent to "Does the cello play as loud as the flute?". Both of these questions correspond to the same functional program even though their text representation is different. Therefore the structures we use include, for each question, a functional representation, and possibly many text representations used to maximize language diversity and minimize the bias in the questions. We have defined 942 such structures.

A template can be instantiated using a large number of combinations of elements. Not all of them generate valid questions. For example "Is the flute louder than the flute?" is invalid because it does not provide enough information to compare the correct sounds regardless of the structure of the scene. Similarly, the question "What is the position of the violin playing after the trumpet?" would be

³<http://sox.sourceforge.net/sox.html>

ill-posed if there are several violins playing after the trumpet. The same question would be considered degenerate if there is only one violin sound in the scene, because it could be answered without taking into account the relation “after the trumpet”. A validation process [11] is responsible for rejecting both ill-posed and degenerate questions during the generation phase.

Thanks to the functional representation we can use the reasoning steps of the questions to analyze the results. This would be difficult if we were only using the text representation without human annotations. If we consider the kind of answer, questions can be organized into 9 families as illustrated in Table 1. For example, the question “What is the third instrument playing?” would translate to the “Query Instrument” family as its function is to retrieve the instrument’s name.

On the other hand we could classify the questions based on the relationships they required to be answered. For example, “What is the instrument after the trumpet sound that is playing the C note?” is still a “query_instrument” question, but compared to the previous example, requires more complex reasoning. The appendix reports and analyzes statistics and properties of the database.

3 Preliminary Experiments

To evaluate our dataset, we performed preliminary experiments with a FiLM network [15]. It is a good candidate as it has been shown to work well on the CLEVR VQA task [11] that shares the same structure of questions as our CLEAR dataset. To represent acoustic scenes in a format compatible with FiLM, we computed spectrograms (log amplitude of the spectrum at regular intervals in time) and treated them as images. Each scene corresponds to a fixed resolution image because we have designed the dataset to include acoustic scenes of the same length in time. The best results were obtained with a training on 35000 scenes and 1400000 questions/answers. It yields a 89.97% accuracy on the test set that comprises 7500 scenes and 300000 questions. For the same test set a classifier choosing always the majority class would obtain as little as 7.6% accuracy.

4 Conclusion

We introduce the new task of acoustic question answering (AQA) as a means to stimulate AI and reasoning research on acoustic scenes. We also propose a paradigm for data generation that is an extension of the CLEVR paradigm: The acoustic scenes are generated by combining a number of elementary sounds, and the corresponding questions and answers are generated based on the properties of those sounds and their mutual relationships. We generated a preliminary dataset comprising 50k acoustic scenes composed of 10 musical instrument sounds, and 2M corresponding questions and answers. We also tested the FiLM model on the preliminary dataset obtaining at best 89.97% accuracy predicting the right answer from the question and the scene. Although these preliminary results are very encouraging, we consider this as a first step in creating datasets that will promote research in acoustic reasoning. The following is a list of limitations that we intend to address in future versions of the dataset.

4.1 Limitations and Future Directions

In order to be able to use models that were designed for VQA, we created acoustic scenes that have the same length in time. This allows us to represent the scenes as images (spectrograms) of fixed resolution. In order to promote models that can handle sounds more naturally, we should release this assumption and create scenes of variable lengths. Another simplifying assumption (somewhat related to the first) is that every scene includes an equal number of elementary sounds. This assumption should also be released in future versions of the dataset. In the current implementation, consecutive sounds follow each other without overlap. In order to implement something similar to occlusions in visual domain, we should let the sounds overlap. The number of instruments is limited to five and all produce sustained notes, although with different sound sources (bow, for cello and violin, reed vibration for the clarinet, fipple for the flute and lips for the trumpet). We should increase the number of instruments and consider percussive and decaying sounds as in drums and piano, or guitar. We also intend to consider other types of sounds (ambient and speech, for example) to increase the generality of the data. Finally, the complexity of the task can always be increased by adding more attributes to the elementary sounds, adding complexity to the questions, or introducing different levels of noise and distortions in the acoustic data.

5 Acknowledgements

We would like to acknowledge the NVIDIA Corporation for donating a number of GPUs, the Google Cloud Platform research credits program for computational resources. Part of this research was financed by the CHIST-ERA IGLU project, the CRSNG and Michael-Smith scholarships, and by the University of Sherbrooke.

References

- [1] Aishwarya Agrawal, Dhruv Batra, and Devi Parikh. “Analyzing the behavior of visual question answering models”. In: *arXiv preprint arXiv:1606.07356* (2016).
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. “Vqa: Visual question answering”. In: *Proc. of ICCV*. 2015, pp. 2425–2433.
- [3] Giuseppe Bandiera, Oriol Romani Picas, Hiroshi Tokuda, Wataru Hariya, Koji Oishi, and Xavier Serra. “Good-sounds.org: a framework to explore goodness in instrumental sounds”. In: *Proc. of 17th ISMIR*. 2016.
- [4] Recommendation ITU-R BS.1770-4. *Algorithms to measure audio programme loudness and true-peak audio level*. Tech. rep. Oct. 2015. URL: https://www.itu.int/dms_pubrec/itu-r/rec/bs/R-REC-BS.1770-4-201510-I!!PDF-E.pdf.
- [5] Jinwei Cao, Jose Antonio Robles-Flores, Dmitri Roussinov, and Jay F Nunamaker. “Automated question answering from lecture videos: NLP vs. pattern matching”. In: *Proc. of Int. Conf. on System Sciences*. IEEE. 2005, 43b–43b.
- [6] Tat-Seng Chua. “Question answering on large news video archive”. In: *Proc. of ISPA*. IEEE. 2003, pp. 289–294.
- [7] Haoyuan Gao, Junhua Mao, Jie Zhou, Zhiheng Huang, Lei Wang, and Wei Xu. “Are you talking to a machine? dataset and methods for multilingual image question”. In: *NIPS*. 2015, pp. 2296–2304.
- [8] Donald Geman, Stuart Geman, Neil Hallonquist, and Laurent Younes. “Visual turing test for computer vision systems”. In: *Proc. of the National Academy of Sciences* 112.12 (2015), pp. 3618–3623.
- [9] Eduard H Hovy, Laurie Gerber, Ulf Hermjakob, Michael Junk, and Chin-Yew Lin. “Question Answering in Webclopedia.” In: *Proc. of TREC*. Vol. 52. 2000, pp. 53–56.
- [10] Mohit Iyyer, Jordan Boyd-Graber, Leonardo Claudino, Richard Socher, and Hal Daumé III. “A neural network for factoid question answering over paragraphs”. In: *Proc. of EMNLP*. 2014, pp. 633–644.
- [11] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. “CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning”. In: *Proc. of CVPR*. IEEE. 2017, pp. 1988–1997.
- [12] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. “Inferring and Executing Programs for Visual Reasoning”. In: *Proc. of ICCV*. Oct. 2017, pp. 3008–3017. DOI: 10.1109/ICCV.2017.325.
- [13] Kyung-Min Kim, Min-Oh Heo, Seong-Ho Choi, and Byoung-Tak Zhang. “DeepStory: video story qa by deep embedded memory networks”. In: *CoRR* (2017). arXiv: 1707.00836.
- [14] “Movieqa: Understanding stories in movies through question-answering”. In: *CVPR*. 2016, pp. 4631–4640.
- [15] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. “Film: Visual reasoning with a general conditioning layer”. In: *CoRR* (2017). arXiv: 1709.07871.
- [16] Deepak Ravichandran and Eduard Hovy. “Learning surface text patterns for a question answering system”. In: *Proc. Ann. Meet. of Ass. for Comp. Ling.* 2002, pp. 41–47.
- [17] Martin M Soubbotin and Sergei M Soubbotin. “Patterns of Potential Answer Expressions as Clues to the Right Answers.” In: *Proc. of TREC*. 2001.
- [18] Christian Steinmetz. *pyLoudNorm*. <https://github.com/csteinmetz1/pyloudnorm/>.
- [19] Ellen M Voorhees et al. “The TREC-8 Question Answering Track Report.” In: *Proc. of TREC*. 1999, pp. 77–82.

- [20] Ellen M Voorhees and Dawn M Tice. “Building a question answering test collection”. In: *Proc. of Ann. Int. Conf. on R&D in Info. Retrieval*. 2000, pp. 200–207.
- [21] Yu-Chieh Wu and Jie-Chi Yang. “A robust passage retrieval algorithm for video question answering”. In: *IEEE Trans. Circuits Syst. Video Technol.* 10 (2008), pp. 1411–1421.
- [22] Hui Yang, Lekha Chaisorn, Yunlong Zhao, Shi-Yong Neo, and Tat-Seng Chua. “VideoQA: question answering on news video”. In: *Proc. of the ACM Int. Conf. on Multimedia*. ACM. 2003, pp. 632–641.
- [23] Peng Zhang, Yash Goyal, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. “Yin and yang: Balancing and answering binary visual questions”. In: *Proc. of CVPR*. IEEE. 2016, pp. 5014–5022.
- [24] Ted Zhang, Dengxin Dai, Tinne Tuytelaars, Marie-Francine Moens, and Luc Van Gool. “Speech-Based Visual Question Answering”. In: *CoRR* abs/1705.00464 (2017). arXiv: 1705.00464. URL: <http://arxiv.org/abs/1705.00464>.
- [25] Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. “Visual7w: Grounded question answering in images”. In: *Proc. of CVPR*. 2016, pp. 4995–5004.

A Statistics on the Data Set

This appendix reports some statistics on the properties of the data set. We have considered the data set comprising 50k scenes and 2M questions and answers to produce the analysis. Figure 2 reports the distribution of the correct answer to each of the 2M questions. Figure 3 and 4 reports the distribution of question types and available template types respectively. The fact that those two distributions are very similar means that the available templates are sampled uniformly when generating the questions.

Finally, Figure 5 shows the distribution of sound attributes in the scenes. It can be seen that most attributes are nearly evenly distributed. In the case of brightness, calculated in terms of spectral centroids, sounds were divided into clearly bright, clearly dark and ambiguous cases (referred to by "None" in the figure). We only instantiated questions about the brightness on the clearly separable cases.

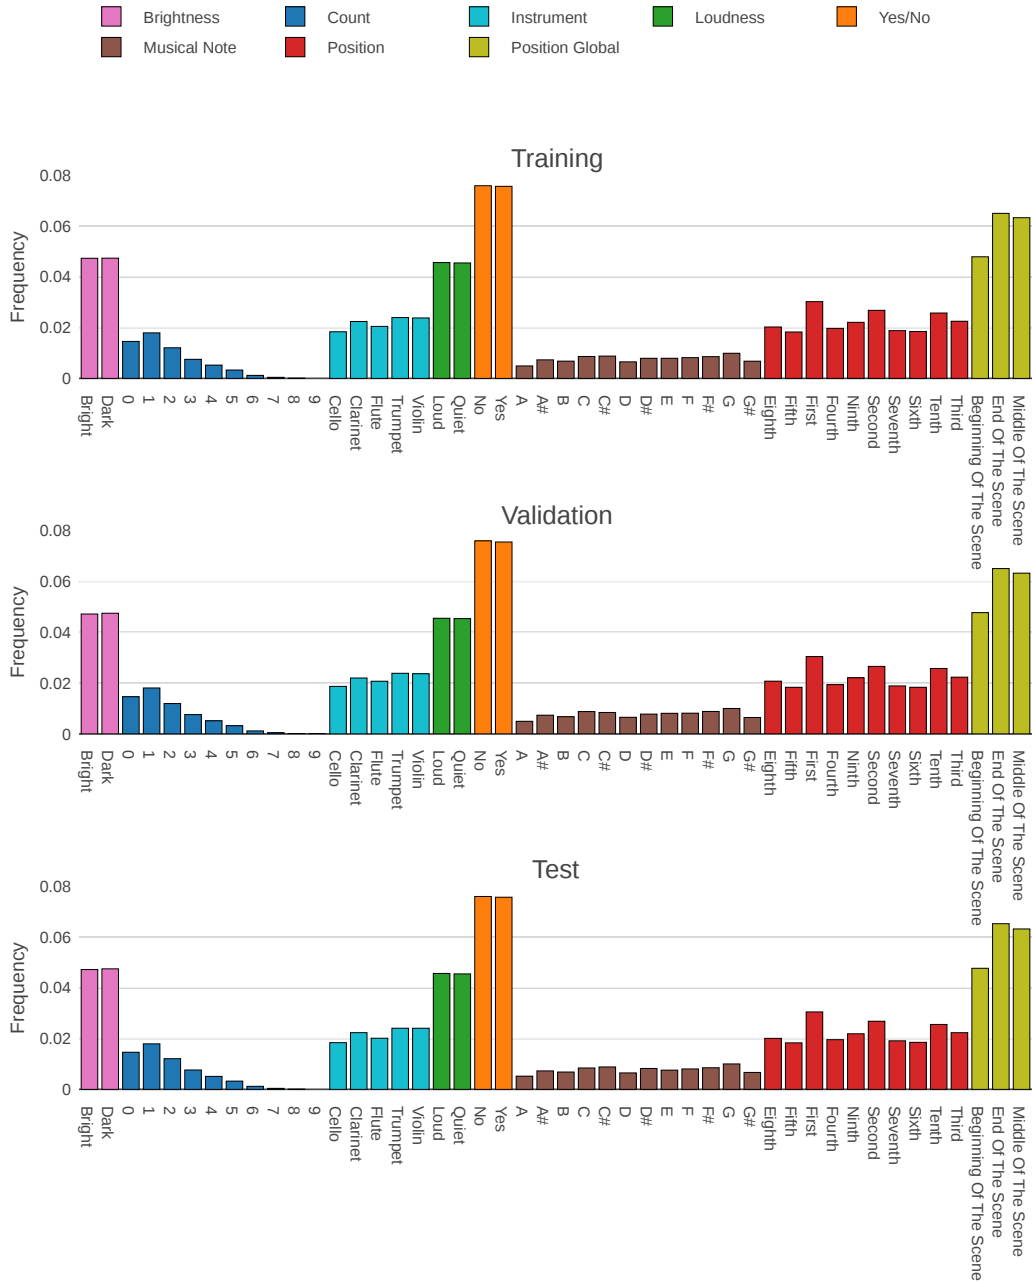


Figure 2: Distribution of answers in the dataset by set type. The color represent the answer category.

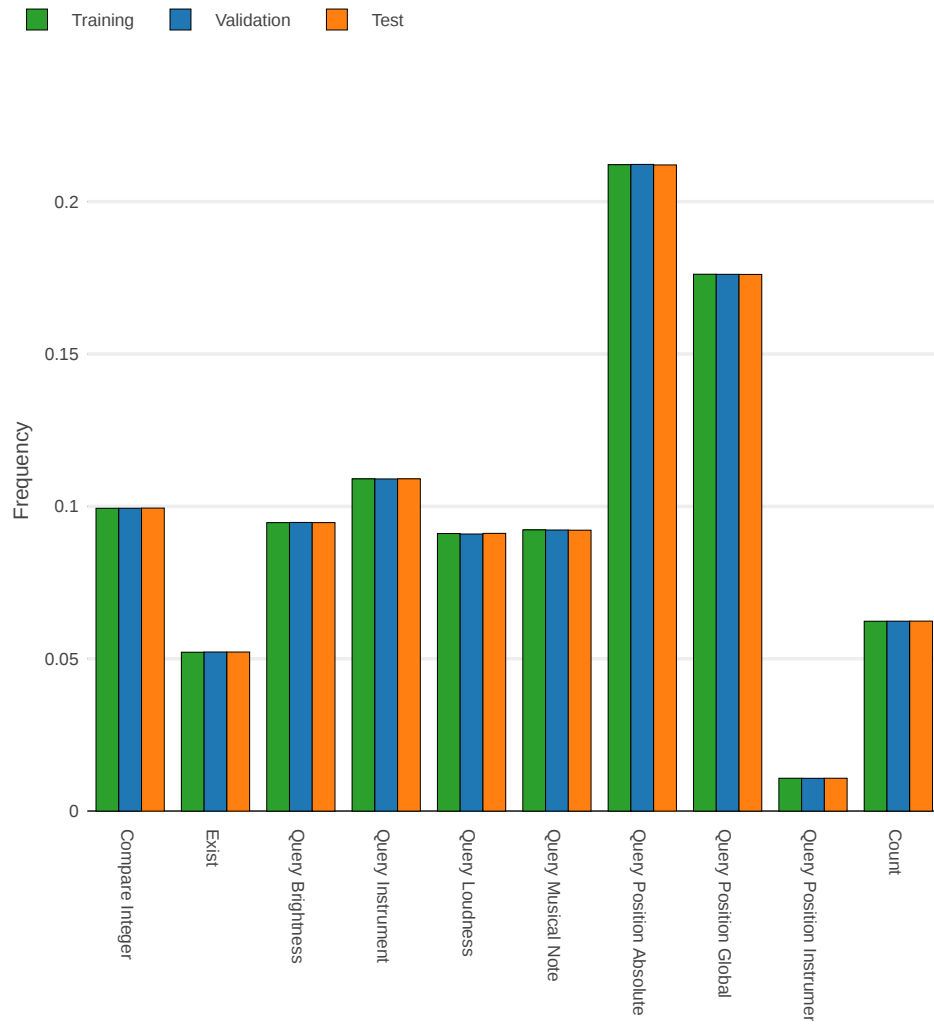


Figure 3: Distribution of question types. The color represent the set type.

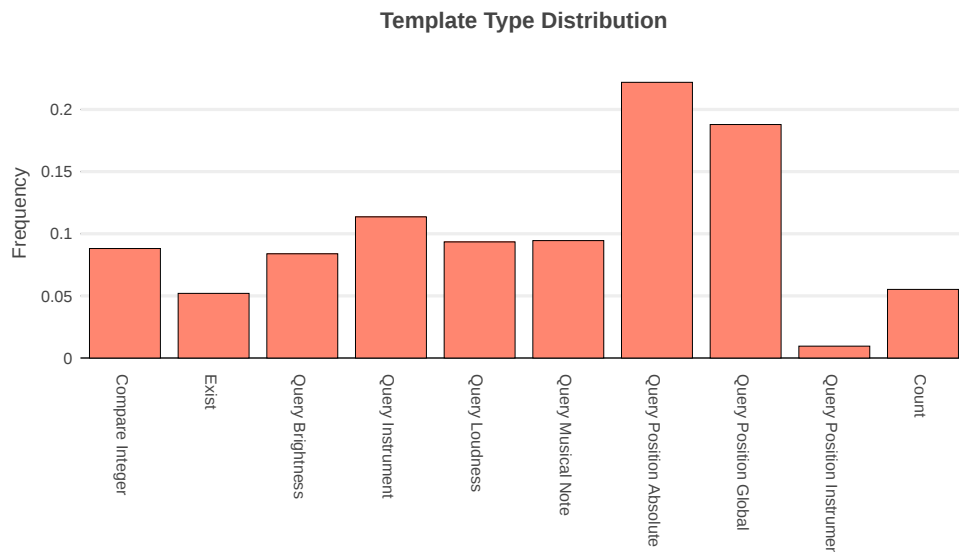


Figure 4: Distribution of template types. The same templates are used to generate the questions and answers for the training, validation and test set.

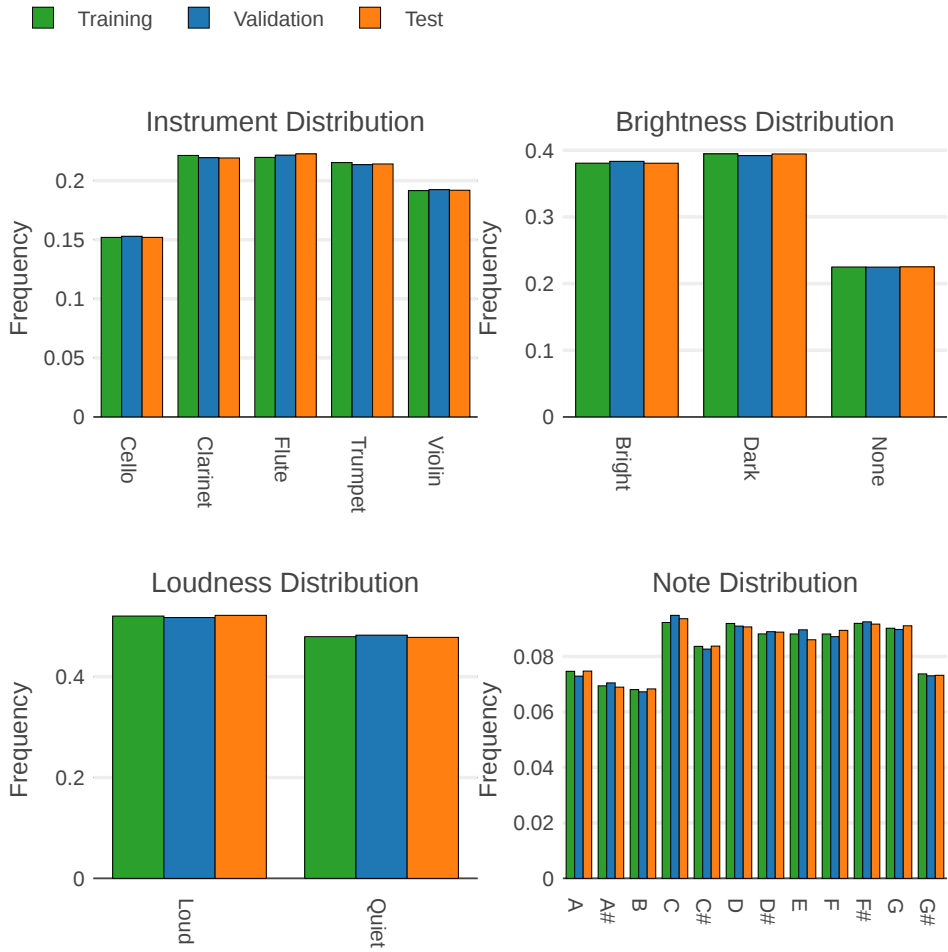


Figure 5: Distribution of sound attributes in the scenes. The color represent the set type. Sounds with a "None" brightness have an ambiguous brightness which couldn't be classified as 'Bright' or 'Dark'.