

# StreamHover: Livestream Transcript Summarization and Annotation

Sangwoo Cho,<sup>†</sup> Franck Deroncourt,<sup>‡</sup> Tim Ganter,<sup>‡</sup> Trung Bui,<sup>‡</sup> Nedim Lipka,<sup>‡</sup>  
Walter Chang,<sup>‡</sup> Hailin Jin,<sup>‡</sup> Jonathan Brandt,<sup>‡</sup> Hassan Foroosh,<sup>†</sup> Fei Liu<sup>†</sup>

<sup>†</sup>University of Central Florida    <sup>‡</sup>Adobe Research

swcho@knights.ucf.edu    {foroosh, feiliu}@cs.ucf.edu  
{deronco, timganter, bui, lipka, wachang, hljin, jbrandt}@adobe.com

## Abstract

With the explosive growth of livestream broadcasting, there is an urgent need for new summarization technology that enables us to create a preview of streamed content and tap into this wealth of knowledge. However, the problem is nontrivial due to the informal nature of spoken language. Further, there has been a shortage of annotated datasets that are necessary for transcript summarization. In this paper, we present StreamHover, a framework for annotating and summarizing livestream transcripts. With a total of over 500 hours of videos annotated with both extractive and abstractive summaries, our benchmark dataset is significantly larger than currently existing annotated corpora. We explore a neural extractive summarization model that leverages vector-quantized variational autoencoder to learn latent vector representations of spoken utterances and identify salient utterances from the transcripts to form summaries. We show that our model generalizes better and improves performance over strong baselines. The results of this study provide an avenue for future research to improve summarization solutions for efficient browsing of livestreams.

## 1 Introduction

One of the most powerful communication mediums is livestreaming. New platforms such as YouTube Live, Twitch, Instagram Live and TikTok encompass a variety of topics, ranging from video games to social media to professional sports. We are particularly interested in livestreams that are distinguished by three characteristics: *Excessive length*, the recordings could last from several minutes to several hours; *Verbal communication*, the use of natural language is the primary means of communication, in contrast to gestures or facial expressions; *Informal nature*, the streamers' language is mostly informal and unplanned, unlike news broadcasts. Without an effective mechanism to summarize such streamed content, livestreaming platforms may not fully meet the needs of their customers.

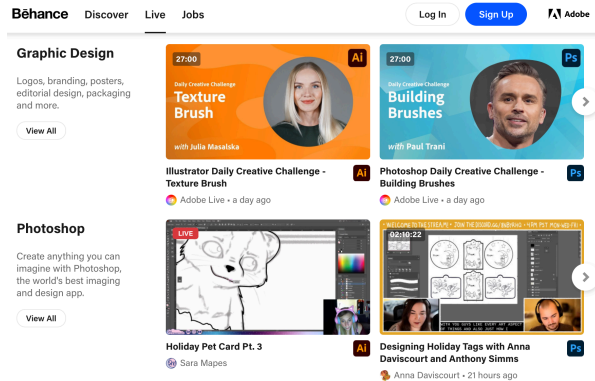


Figure 1: An example of streamed content on Bēhance, a streaming platform for artists and designers to showcase creative work related to Adobe Photoshop, Illustrator, Fresco, UI/UX, photography and more. TOP: The videos are each 27 minutes long. BOTTOM: One video is being broadcast live, the other is >2 hours long.

Our goal in this work is to create a *text preview* of the streamed content. When a user **hovers over** the thumbnail or scrolls past a video, they are shown a preview of the content. We present a dataset of over 500 hours of video footage, which were streamed live on a social media platform (behance.net) created to showcase and discover creative work. Figure 1 shows an example of the streams, where the artists showcase the use of Adobe Photoshop and Illustrator in designing holiday cards and posters. It is necessary to point out that video analysis is not suitable here, as the video only mirrors the artists' screen content. As a first step towards automatic creation of a text preview, we focus on identifying salient utterances to produce an extract from the livestream transcript.

We make use of vector-quantized variational autoencoders (VQ-VAE; van den Oord et al., 2017) to identify salient utterances. The model has been applied successfully to opinion summarization that learns in-domain sentence representations (Angelidis et al., 2021), which is essential for adaptation of general-domain models. We refrain from using sequential methods for utterance selection. First,

it is difficult to scale up sequential prediction to process transcripts that exceed the maximum allowed length, even with models that handle long text (Beltagy et al., 2020; Zhao et al., 2020). Second, sequential methods (Narayan et al., 2018b; Xiao and Carenini, 2019) may not give enough flexibility to select salient utterances *on-the-fly* when content is being streamed live, thus they are unsuitable for our case.

There has been a shortage of annotated datasets that are necessary for livestream transcript summarization. We build a browser-based user interface for summary annotation that provides to the annotators a clip of the livestream recording alongside a synchronized display of the transcript. The interface allows annotators to conveniently label summary utterances and write an abstractive summary using their own words (Figure 3). With a total of 500 hours of annotated video footage, our dataset is notably larger than existing annotated corpora for transcript summarization (Janin et al., 2003; Carletta et al., 2006). We compare our summarization approach with strong baselines on the dataset and shed light on the task of livestream transcript summarization. Our contributions are as follows.

- We create a detailed annotation interface and new benchmark dataset for automatic summarization of livestream transcripts. An informative preview of streamed content is of crucial importance to users when considering whether to hit play.
- We present **StreamHover**, a unsupervised model based on VQ-VAE to identify salient utterances from livestream transcripts to form preview summaries. We evaluate the method across multiple dimensions and discuss its strengths and weaknesses. Empirical results show that our method outperforms strong summarization baselines.<sup>1</sup>

## 2 Related Work

Closed captions are often provided onscreen, turning streaming videos into text on an unprecedented scale (Besik, 2020). However, there are very few summarization studies that attempt to generate text previews of streaming videos to help users browse or refind information that has been watched before. Neural text summarizers have focused primarily on written text, including news articles, reviews, scientific papers and book chapters (See et al., 2017; Tan

<sup>1</sup>Our annotations and source code are available at <https://github.com/ucfnlp/streamhover>

07:35

OK. It looks like we are live. How is everyone today good morning Good afternoon good evening? Wherever you are in the world my name is Jason Levine and for today's brief live stream go to be talking about a topic which to be honest. I didn't think I'd revisit again seeing as I've got some archive videos on this. But in the last week alone. I had probably 3 or 4 different customers on YouTube and a couple more on Twitter and even if you on Facebook asked me how the heck do you burn an audio CD in Adobe audition?

08:10

And specifically what else do you need to do to make sure that you know everything is right when you're burning it to CD so that's what I thought I would take the opportunity to do today, so as always. We are coming to you live on Facebook YouTube be hands and Twitter Periscope, so we've got our live interactive chat. Thank you so much for joining as always, I will of course, be checking the chats at the end of the broadcast, but we got it just really it's a couple of quick things to cover here? What is up Catherine Lacoste. An it's a really simple process but there's definitely a few things that you need to be aware of the need to know.

Figure 2: A transcript snippet from “How to Create an Audio CD in Adobe Audition.” Most utterances are off-topic, except for the one marked in blue, suggesting the information density of livestream transcripts is low.

et al., 2017; Chen and Bansal, 2018; Narayan et al., 2018a; Gehrmann et al., 2018; Cohan et al., 2018; Liu and Lapata, 2019; Fabbri et al., 2019; Bražinskas et al., 2020; Ladhak et al., 2020; Song et al., 2021). Despite their success, it remains unclear as to if and how the summarizers can be extended to spoken text, whose utterances may have very low information density.

It is crucial to identify salient content from transcripts where a substantial number of utterances are devoted to informal chit-chats in an attempt to connect with the audience (Figure 2). We investigate extractive rather than abstractive approaches as the latter are prone to generate hallucinated content that does not exist in the source text (Cao et al., 2017; Kryscinski et al., 2019; Lebanoff et al., 2019; Maynez et al., 2020). The problem could be exacerbated by ungrammatical spoken utterances and transcription errors. Instead, we consider VQ-VAE, an unsupervised representation learning technique (van den Oord et al., 2017; Jin et al., 2020; Angelidis et al., 2021) for content extraction. Unsupervised training of the VQ-VAE model and its inference could potentially be performed at the same time, allowing important utterances to be extracted from a transcript segment *on-the-fly* during streaming, without interrupting the learning process. It is also easier to tailor the model to specific domains compared to contemporary extractive methods (Yasunaga et al., 2017; Dong et al., 2018; Xu and Durrett, 2019; Wang et al., 2020).

Our work contributes to a refined understanding of transcript summarization, which is understudied relative to its importance and potential. The transcripts may be obtained from channels such as movies and TVs (Papalampidi et al., 2020; Chen et al., 2021), interviews (Zhu et al., 2021), multi-party meetings (Murray and Carenini, 2008; Wang and Cardie, 2013; Li et al., 2019b; Koay et al., 2020, 2021; Zhong et al., 2021), telephone speech (Kafle and Huenerfauth, 2018) and more. The main thrust distinguishing our work with others is the combination of a benchmark summarization dataset, novel summarization methods and a challenging new domain where salient content is scattered throughout the transcript and mixed with substantial chit-chats. We do not make use of video event detection or multi-modal fusion (Zhu et al., 2018; Palaskar et al., 2019; Li et al., 2020) as little information could be gleaned from videos that mirror the artists’ desktop. Instead, we focus on generating short descriptions from transcripts and leave for future work cross-modality research. We describe our data annotation process in the following section.

### 3 Our Dataset

We aim to create a large and representative corpus containing transcripts and summaries of streamed videos. We explore a leading social media platform (Behance.net) supported by Adobe Creative Cloud that features livestreams of creative work by artists and designers. The website boasts over 10 million users, who watch artists and designers livestream when they create. Our data are extracted from this website, containing a large quantity of streamed videos (>5,000), the length of which ranges from minutes to several hours. The streamers’ language is unplanned, instead of rehearsed as that of TED talks (Hernandez et al., 2018).

We obtain a total of 5,398 streamed videos. The metadata of a video includes its ID, duration, title, a short description and the transcript. Automatic transcription was provided by Microsoft Automatic Speech Recognition which helps make videos accessible to a wide audience. Each transcript contains a set of segments, each corresponds to about 30 seconds of audio. Each segment contains a set of utterances.<sup>2</sup> Figure 2 shows an example of the segments and utterances. The offset of the segment

<sup>2</sup>An utterance is the smallest unit of speech in improvised spoken language. It usually begins and ends with a clear pause. We obtain utterance boundaries from Microsoft’s ASR system.

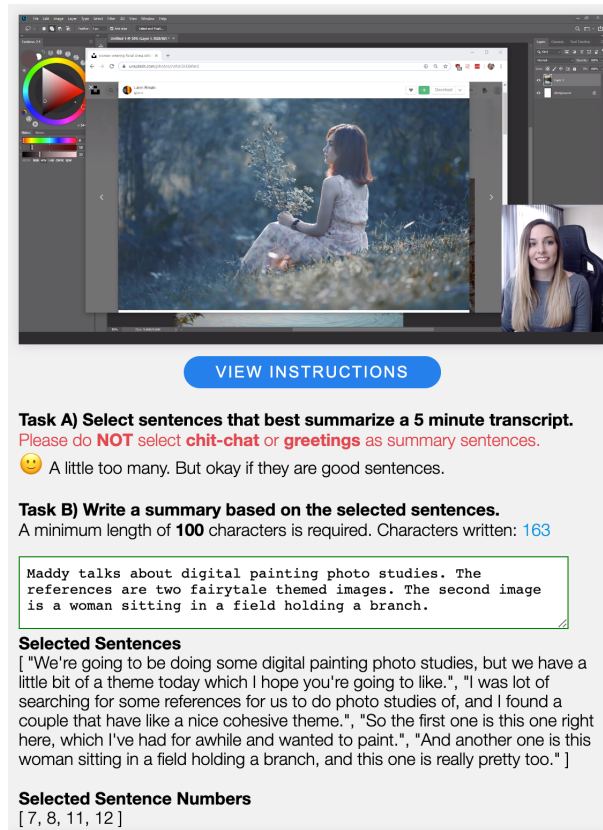


Figure 3: An example of our browser-based annotation interface. It includes a clip of the streamed video alongside a display of the transcript (omitted for space). The streamer talks about *Digital Painting with Maddy Bellwoar* to create fairytale themed images. The annotators are asked to write a concise summary of this clip using their own words (Task A) and identify summary utterances (Task B).

indicates the number of minutes since the beginning of the recording.

When a user hovers over the thumbnail or scrolls past a video, we expect a textual summary to give a glimpse of the verbal content. This view of summarization leads us to annotate salient content across the video in an *equally detailed* manner. It naturally avoids lead bias that is ubiquitous in news (Grenander et al., 2019). We segment a video into 5-minute clips and annotate each clip for summary-worthy content. A clip contains an average of 51 utterances and 460 words. Due to time and budget constraints, we select 370 streamed video for summary annotation.<sup>3</sup> Table 1 provides a detailed comparison of our annotated corpus with previous datasets, including Switchboard (Godfrey et al., 1992), ICSI (Janin et al., 2003) and AMI (Carletta et al., 2006) that contain both transcripts and human-annotated ex-

<sup>3</sup>Details of video selection are provided in Supplementary.

Dataset	Type	Duration	Extr.	Abst.
Switchboard	Telephone	300 Hrs	No <sup>†</sup>	No
ICSI	Meeting	75 Hrs	Yes	Yes
AMI	Meeting	100 Hrs	Yes	Yes
<b>StreamHover</b>	<b>Livestream</b>	<b>500 Hrs</b>	<b>Yes</b>	<b>Yes</b>

Table 1: A comparison of the transcript summarization datasets with manually annotated extractive/abstractive summaries. “Yes/No” indicate a summary type is available or not. † suggests only small pilot summary annotations are available for the Switchboard dataset (Penn and Zhu, 2008). With a total duration of over 500 hours, our dataset is notably larger than similar datasets.

tractive/abstractive summaries. With a combined duration of 500 hours, our dataset is substantially larger than previously released datasets.

We recruit 12 workers from Upwork.com and validate their language skills for summary annotation. Upwork is a freelancing platform that allows us to reach out to workers directly to ensure our instructions are fully understood. Each worker is asked to write a concise summary for a given clip using their own words (Task A) and identify summary utterances (Task B) using the graphical user interface (Figure 3), which shows a clip of the streamed video alongside a synchronized display of the transcript. Additionally, our guidelines suggest a good summary of Task A should have at least 100 characters and that of Task B should have between 50 and 80 words ( $\sim 15\%$  compression). As is the case with meeting summaries (Janin et al., 2003), a clip is annotated by a single worker owing to an expensive annotation process. The worker can also identify a clip to be chitchat, in which case it will not be annotated for summaries.

Table 2 shows our dataset statistics. On average, a human *abstract* contains 3 sentences (36 words) and a human annotated *extract* contains 5.5 utterances (80 words). Moreover, summary utterances constitute 8.9% and 8.6% of all utterances in terms of number and duration. We study inter-annotator agreement by inviting three workers to each annotate 8 hours of video that contains a total of 96 clips. Using 10-second intervals as measuring units,<sup>4</sup> the Fleiss’ Kappa on identifying summary utterances is 0.13. We note that the score is within the range of what is normally found in annotating speech transcripts for extractive summaries (0.1 $\sim$ 0.3; Marge

<sup>4</sup>We use 10-second intervals rather than utterances as measuring units as the duration of utterances vary. If annotators all selected some content, or no content at all, from a 10-second interval, they are in agreement.

The Bēhance Corpus	
Total number of annotated videos	370
Total annotated duration in hours	500
Total number of utterances	331,928
Average number of utterances in a clip	61.23
Average duration of utterances in seconds	3.04
Average number of words in an utterance	9.48
Total number of annotated clips	6,003
Total number of chitchat clips	582
Total number of human annotators	12
Avg. # of (sentences) words in a human abstract	(3.0) 36
Avg. # of (utterances) words in a human extract	(5.5) 80
Percentage of duration of summary utterances	8.57%

Table 2: Statistics of our dataset.

et al., 2010), as annotating spoken text is a highly challenging task. We find that annotators tend to perceive the same region as salient but they may disagree as to which utterances should be included in the summary due to verbosity of spoken text. We refer the reader to (Artstein and Poesio, 2008) for interpretations and improvements to IAA.

## 4 Summarization

Let  $\mathcal{X}$  denote a sequence of spoken utterances from a segment of the transcript. Our summarizer aims to extract a subset of utterances  $\mathcal{Y} \subset \mathcal{X}$  that convey the essential content of the input. We experiment with an unsupervised summarizer that leverages vector-quantized variational autoencoders (VQ-VAE; van den Oord et al., 2017) to learn utterance representations and identifies summary utterances. The method was explored for opinion summarization (Angelidis et al., 2021) and machine translation (Prato et al., 2020). We are interested in using the method to account for domain characteristics of livestreams, which showcase new and creative work of artists and designers on their use of Photoshop, Illustrator, and other tools.<sup>5</sup>

VQ-VAE is a powerful framework for learning latent variable models using deep neural networks. It learns *discrete* vector representations for an utterance, which is then used to categorize the utterance along various dimensions. E.g., “*Good morning Hi Everybody*” suggests a greeting and opens up a dialogue; “*I had probably 3 or 4 different customers on YouTube and ... on Facebook asked me how the heck do you burn an audio CD in Adobe Audition*” engages the audience and introduces the

<sup>5</sup>A show may have more than one host, their utterances are treated indistinguishably due to speaker diarization that identifies different speakers in the audio is not provided.




Figure 4: Our summarizer embeds an input utterance using BERT, transforms BERT’s semantic space to a set of latent codes, then reconstructs the utterance using the code embeddings. We identify summary utterances as those associated with prominent latent codes/topics. The model is trained using a dictionary learning algorithm for code embeddings ( $\mathbf{E}$ ) and backpropagation with a straight-through estimator for model parameters  $\theta$ ,  $\varphi$  and  $\phi$ .

main topic. The VQ-VAE method groups utterances based on their discrete representations and selects salient utterances to form a summary.

We employ an embedding function  $\text{Embed}_\theta(\cdot)$  to map an input utterance  $x$  into a semantically meaningful space. The space is subsequently discretized according to a codebook. To achieve this, we prefix  $x$  with a [CLS] token and append a [SEP] token, pass it into a BERT model, then obtain the vector corresponding to [CLS] as a pooled representation of the utterance, denoted by  $\mathbf{h} \in \mathbb{R}^H$  (Eq. (1)). We use a  $\text{ConvEncoder}_\varphi(\cdot)$  with a set of  $D$  filters to convolve the input  $\mathbf{h}$ . The output is a sequence of feature vectors  $[\mathbf{q}_1, \dots, \mathbf{q}_H]$  where  $\mathbf{q}_i \in \mathbb{R}^D$  (Eq. (2)). We define a codebook  $\mathbf{E} = [\mathbf{e}_1, \dots, \mathbf{e}_K]$ , where  $K$  is the number of latent codes and  $\mathbf{e}_k \in \mathbb{R}^D$  is the  $k$ -th code embedding. The  $i$ -th feature  $\mathbf{q}_i$  is assigned to the latent code  $z_i$  whose embedding  $\mathbf{e}_{z_i}$  has the minimum Euclidean distance with it (Eq. (3)). Our method essentially discretizes the  $H$ -dimensional semantic space by producing latent codes  $\{z_i\}_{i=1}^H$ , one for each dimension of the semantic space.

$$\mathbf{h} = \text{Embed}_\theta(x) \in \mathbb{R}^H \quad (1)$$

$$[\mathbf{q}_1, \dots, \mathbf{q}_H] = \text{ConvEncoder}_\varphi(\mathbf{h}), \mathbf{q}_i \in \mathbb{R}^D \quad (2)$$

$$z_i = \arg \max_k -\|\mathbf{q}_i - \mathbf{e}_k\|_2, i \in [H] \quad (3)$$

With the latent code embeddings  $[\mathbf{e}_{z_1}, \dots, \mathbf{e}_{z_H}]$ , we seek to reconstruct the input utterance, which is achieved by generating a dense vector  $\tilde{\mathbf{h}}$  using a  $\text{ConvDecoder}_\varphi(\cdot)$  (Eq. (4)).  $\tilde{\mathbf{h}}$  is then fed to a Transformer decoder to reconstruct the original utterance  $\tilde{x}$  (Eq. (5)). In this process, the code embeddings serve as “topic vectors” that group dimensions of the semantic space into clusters relevant to the application domain. Our model parameters include those used by the BERT encoder and Transformer decoder ( $\theta$  and  $\phi$ ), the convolutional encoder and

decoder that use tied parameters ( $\varphi$ ), and embeddings of the codebook  $\mathbf{E}$ .

$$\tilde{\mathbf{h}} = \text{ConvDecoder}_\varphi([\mathbf{e}_{z_1}, \dots, \mathbf{e}_{z_H}]) \in \mathbb{R}^H \quad (4)$$

$$\tilde{x} = \text{Generate}_\phi(\tilde{\mathbf{h}}) \quad (5)$$

We next describe the loss function used to learn these parameters. The loss function of our model comprises of three parts, including a cross-entropy term between the original and reconstructed utterance  $\text{xEnt}(x, \tilde{x})$  that optimizes the BERT embedder  $\theta$ , Transformer generator  $\phi$ , and convolutional encoder and decoder  $\varphi$ , as shown in Figure 4. The gradients will, however, bypass the latent code embeddings due to the straight-through estimator (Benigno et al., 2013). To learn code embeddings in an end-to-end manner, we use a dictionary learning algorithm (van den Oord et al., 2017) that moves code embeddings  $\mathbf{e}_{z_i}$  towards feature vectors  $\mathbf{q}_i$  by minimizing the  $l_2$ -distance between the two vectors  $\|\mathbf{e}_{z_i} - \text{sg}(\mathbf{q}_i)\|_2^2$ , where  $\text{sg}(\cdot)$  is a stop-gradient operator that constrains its operand to be a non-updated constant during backpropagation, i.e., it stops  $\mathbf{q}_i$  from being updated. As illustrated in Eq. (6), we additionally apply a commitment loss to encourage the feature vector  $\mathbf{q}_i$  to commit to a code embedding.  $\|\text{sg}(\mathbf{e}_{z_i}) - \mathbf{q}_i\|_2^2$  prevents  $\mathbf{q}_i$  from deviating too much from the code embedding  $\mathbf{e}_{z_i}$ . This loss term is associated with a coefficient  $\beta \in [0, 1]$ .

$$\mathcal{L}(\theta) = \text{xEnt}(x, \tilde{x}) + \sum_i \|\mathbf{e}_{z_i} - \text{sg}(\mathbf{q}_i)\|_2^2 + \beta \sum_i \|\text{sg}(\mathbf{e}_{z_i}) - \mathbf{q}_i\|_2^2 \quad (6)$$

At test time, we define summary utterances as those associated with prominent latent codes/topics. Given a set of  $N$  utterances, we obtain latent codes from the  $n$ -th utterance using Eq. (3), denoted by  $\{z_i^{(n)}\}_{i=1}^H$ . This gives a total of  $N \times H$  codes from which we find prominent ones. They are denoted by

$\mathcal{P}$  which contains a set of most frequently occurring codes. A score  $\mathcal{S}(x_n)$  is assigned to utterance  $x_n$  that computes how often it is associated with those prominent codes  $\mathcal{P}$ . In Eq. (7),  $\sum_{i=1}^H \mathbb{1}[z_i^{(n)} = k]$  indicates the number of times the  $n$ -th utterance is assigned to code  $k$ , where  $k$  belongs to  $\mathcal{P}$ . Finally, we extract  $K$  highest-scoring utterances to form an extractive summary of the input.

$$\mathcal{S}(x_n) = \sum_{k \in \mathcal{P}} \sum_{i=1}^H \mathbb{1}[z_i^{(n)} = k] \quad (7)$$

Our method draws on the convolutional encoder and decoder to transform BERT’s semantic space to map each dimension to a latent code. The summary selection process is deterministic and our encoder takes full advantage of a large, pre-trained model to produce initial utterance representations. This design sets our method apart from that of Angelidis et al. (2021). Moreover, the method has the potential for modelling topic transitions between utterances to improve summarization of livestreams, which we leave for future work.

## 5 Experiments

**Dataset.** Finding salient content from livestream transcripts is a “needle-in-the-haystack” problem. Our summarization dataset contains a total of 370 videos split into short clips of 5 minutes each. The annotators manually annotated 5,421 clips ( $\sim 451$  hours) with extractive and abstractive summaries. 582 clips ( $\sim 49$  hours) are removed because they are identified to contain only chit-chats. The dataset is divided into training, validation and test splits:

- 3,884 clips (320 videos / 323 hours) in training,
- 728 clips (25 videos / 61 hours) in validation,
- 809 clips (25 videos / 67 hours) in test split.

We call our summarizer “**StreamHover**.” When a user hovers their mouse over a video’s timeline, a summary preview is shown and keeps updating. As a first attempt, StreamHover focuses on extracting salient utterances from individual clips instead of whole streams to encourage selected utterances to be mostly evenly distributed across the stream. When the content is provided live, the stream can be divided into short clips and our algorithm consumes one clip at a time to produce summaries on-the-fly. It is important to note that extracting summary utterances remains challenging even for modern neural summarizers. E.g., Kedzie et al. (2018) reveal that summarizers may not effectively identify summary content without a dependency on intentional

lead bias in news writing. Our setting is challenging as not only are there few utterances deemed to be summary-worthy but such utterances can occur anywhere in a video clip.

**Baselines.** We compare StreamHover with state-of-the-art extractive and abstractive summarizers. The abstractive summarizers generate an abstract from the transcript of a clip without tuning.<sup>6</sup> These include BART-large, BART-large-cnn (Lewis et al., 2020) and T5 (Raffel et al., 2020), which are some of the strongest performing neural abstractive summarizers that are pre-trained on language modeling and summarization tasks.

The unsupervised extractive summarizers extract salient utterances from a clip. LexRank (Erkan and Radev, 2004) and TextRank (Mihalcea and Tarau, 2004) are graph-based models that extract relevant sentences based on eigenvector centrality. Sum-Basic (Vanderwende et al., 2007) assigns higher scores to sentences containing frequently occurring content words. We further compare to a novel unsupervised graph-based summarization method for speech transcripts: FluCovRank (Shang et al., 2018) groups utterances into clusters, generates an abstractive sentence from each cluster, then selects the best elements from abstractive sentences under a budget constraint. Finally, we compare our approach with the Quantized Transformer (Angelidis et al., 2021), which uses a clustering interpretation of the quantized space and two-step sampling algorithm to extract summary sentences from reviews.

**Settings.** We use pretrained BERT-BASE as our embedder  $\text{Embed}_\theta(\cdot)$ . The model has 12 layers, 12 heads per layer and a hidden size ( $H$ ) of 768. A 6-layer Transformer decoder is used as the generator  $\text{Generate}_\phi(\cdot)$  to reconstruct the original utterance. The model has 8 heads per layer, a hidden size of 768, and randomly initialized parameters. The convolutional encoder and decoder use a kernel size of 3. Because our embedder is pretrained and the remaining parameters are not, we divide them into two groups  $\mathcal{E}=\{\theta\}$  and  $\mathcal{R}=\{\phi, \varphi\}$ , then apply separate training schedules. Following Liu and Lapata (2019) we use two Adam optimizers:

$$\begin{aligned} lr_{\mathcal{E}} &= \tilde{lr}_{\mathcal{E}} \cdot \min(\text{step}^{-0.5}, \text{step} \cdot \text{warmup}_{\mathcal{E}}^{-1.5}), \\ lr_{\mathcal{R}} &= \tilde{lr}_{\mathcal{R}} \cdot \min(\text{step}^{-0.5}, \text{step} \cdot \text{warmup}_{\mathcal{R}}^{-1.5}) \end{aligned}$$

where the learning rate for the embedder  $\tilde{lr}_{\mathcal{E}}=7e^{-4}$

<sup>6</sup>In a similar vein, our summarizer uses the transcripts to learn model parameters. It does not require utterance labels.

System	3-Sentence Output				4-Sentence Output				5-Sentence Output			
	P (%)	R (%)	F (%)	#Wrds	P (%)	R (%)	F (%)	#Wrds	P (%)	R (%)	F (%)	#Wrds
LEAD-N	18.83	9.57	12.5	38.53	18.63	12.61	14.77	51.35	18.71	15.76	16.82	64.04
SumBasic	8.32	4.15	5.45	29.44	8.47	5.61	6.63	39.97	8.83	7.44	7.92	51.54
QuantizedTran	10.40	13.35	11.07	80.09	10.44	17.67	12.60	104.66	10.58	21.86	13.72	128.35
LexRank	23.94	12.14	15.86	59.51	23.34	15.96	18.57	77.33	23.47	20.03	21.19	94.43
TextRank	30.45	15.37	20.10	73.35	28.18	18.92	22.24	92.46	27.00	22.59	24.17	110.42
<b>StreamHover</b>	<b>36.18</b>	<b>18.21</b>	<b>23.87</b>	<b>88.04</b>	<b>34.86</b>	<b>23.29</b>	<b>27.52</b>	<b>113.40</b>	<b>33.92</b>	<b>28.42</b>	<b>30.47</b>	<b>137.02</b>

Table 3: Classification performance of extractive summarizers on identifying ground-truth summary utterances.

	System	R-1	R-2	R-L	#Wrds
Abstract	BART-large	22.98	8.35	15.04	123.05
	BART-large-cnn	23.03	8.03	16.62	43.13
	T5-large	24.20	8.55	17.56	50.98
	FluCovRank	25.29	10.93	16.28	50.00
Extract	LEAD-5	24.65	9.54	17.59	64.04
	SumBasic	23.15	5.57	14.76	51.54
	QuantizedTran	23.90	7.90	15.37	128.35
	LexRank	<b>26.14</b>	10.18	18.24	94.43
	TextRank	25.94	11.44	18.89	110.42
	<b>StreamHover</b>	<b>25.62</b>	<b>12.70</b>	<b>19.33</b>	<b>137.02</b>
	Oracle (Extract)	43.42	30.58	37.99	110.51

Table 4: Results of extractive and abstractive summarizers evaluated by ROUGE. Extractive summarizers generate a 5-utterance summary for each clip. *Oracle* contains ground-truth summary utterances. StreamHover achieves the highest scores on R-2 and R-L.

	System	Fluency	Informa.	Overall
Human	FluCovRank	-0.95	-0.93	-0.97
	LexRank	0.25	0.11	0.17
	BART-large	0.28	0.31	0.28
	<b>StreamHover</b>	<b>0.42</b>	<b>0.52</b>	<b>0.52</b>

Table 5: Results of human evaluation regarding fluency, informativeness and the overall quality of system summaries using Best-Worst Scaling.

is smaller than that of the rest params  $\tilde{l}r_{\mathcal{R}}=4e^{-2}$ . Its warmup period is longer:  $warmup_{\epsilon}=3,000$  for the embedder and  $warmup_{\mathcal{R}}=1,500$  for the rest. It allows the pretrained embedder to be updated in a slower pace until other model parameters start to generate accurate gradients.

All of our models are trained for 30 epochs on dual NVIDIA V100 GPUs with gradient accumulation every ten steps. We experiment with different numbers of filters,  $D = \{64, 100, 128\}$ , for the convolutional encoder and decoder. The number of latent codes are varied in  $K = \{512, 1024, 2048\}$ . The coefficient  $\beta$  used for commitment loss is set to 0.25 (Eq. (6)). These hyperparameters are tuned on the validation set. We keep only utterances that contain  $>5$  words in consideration. The final training set contains 168,111 utterances.




Figure 5: The proportion of times a system is selected as the “Best.” FluCovRank is omitted as it is  $<1\%$ .

## 5.1 Results

In Table 3, we analyze the performance of extractive summarizers on identifying ground-truth summary utterances and report their precision, recall and  $F_1$ -measure scores. We vary the length of their output to yield  $\{3, 4, 5\}$ -utterance summaries. In comparison, a ground-truth extract contains 5.5 utterances. The Lead-N baseline selects the first N utterances of a clip. It gives low scores because our data do not present strong lead bias as that of news articles. We find that StreamHover consistently outperforms other summarization systems across all lengths. Its length, when measured by number of words, is comparable to that of LexRank and TextRank. The highest  $F_1$ -score of 30.47% is achieved when StreamHover generates a 5-utterance summary for each 5-minute clip. This amounts to rendering one utterance per one-minute segment when a user scrolls past the video.

In Table 4, we compare extractive and abstractive summarizers and report ROUGE scores (Lin, 2004) that measure content overlap between system and reference summaries.<sup>7</sup> We use human abstracts as the reference. All extractive summarizers produce 5-utterance summaries and *Oracle Extract* contains ground-truth utterances. It places an upper bound on the performance of extractive summarizers. We observe that StreamHover yields the highest scores on both R-2 and R-L metrics.

<sup>7</sup>The recent automatic metrics (Zhang et al., 2020; Sellam et al., 2020) have not been tested on speech transcripts. Spoken text contains filled pauses (um, uh, well), disfluencies (go-go-go away), repetitions and verbal interruptions. ROUGE is the only metric that has been validated to attain good correlation with human judgments on transcripts (Liu and Liu, 2008).

<p><b>FluCovRank</b></p> <ul style="list-style-type: none"> <li>• top left bottom / cloud studies today / find links to their original posts / hey jennifer saw the images / love the top left and bottom / info tab and i uploaded / colors are beautiful but im partial through colorful sky scenes / pretty large about 4000 by 4000 pixels / photo studies of today / moment</li> </ul>
<p><b>LexRank</b></p> <ul style="list-style-type: none"> <li>• I hope you guys are having a good day so far.</li> <li>• So I'm going to be painting from these images and these beautiful photos are from various photographers.</li> <li>• Those yeah well top right also is like very Contra high contrast that tends to like grab my attention when I look at the sheet but I would say <u>top left and bottom right give me the most like happy feels.</u></li> <li>• So yeah, if you guys want to grab the reference images, you can find them in the stream description below the individual images...</li> </ul>
<p><b>BART-Large</b></p> <ul style="list-style-type: none"> <li>• Hello good morning everybody welcome to the stream. I hope you guys are having a good day so far. Is there a lot of buffering or are we doing alright? I got a little message that there was some connectivity issue. For a moment there, so I hope I hope it's OK. Yeah, I'll just keep going. So yeah, if you guys want to grab the reference images, you can find them in the stream description below the individual images...</li> </ul>
<p><b>Quantized Transformer</b></p> <ul style="list-style-type: none"> <li>• Good to see you were going to be doing cloud studies today.</li> <li>• The stream in the description.</li> <li>• One of them is from Morguefile, One is from unsplash, well, <u>two are from Unsplash and one is from pixels</u> there a little bit from all over the place, but you can find the photographers below if you'd like.</li> <li>• Hey Jennifer, saw the images.</li> <li>• Let's see top left, bottom right...</li> </ul>
<p><b>StreamHover (Ours)</b></p> <ul style="list-style-type: none"> <li>• So if anybody is interested in joining in, if you want to <u>work on some skies for your landscapes for future landscapes, this is what we're going to be doing.</u></li> <li>• One of them is from Morguefile, One is from unsplash, well, <u>two are from Unsplash and one is from pixels</u> there a little bit from all over the place, but you can find the photographers below if you'd like.</li> <li>• Those yeah well top right also is like very Contra high contrast that tends to like grab my attention when I look at the sheet but I would say <u>top left and bottom right give me the most like happy feels.</u></li> <li>• So yeah, if you guys want to grab the reference images, you can find them in the stream description below the individual images...</li> </ul>

Table 6: Example system summaries for *Digital Painting Studies with Maddy Bellwoar-Clouds*. The BART summary is fluent but its content lacks specificity, as is the case for LexRank. The summary segments selected by FluCovRank are ungrammatical. StreamHover identifies on-topic and informative utterances related to digital painting. Their relevant spans of text are manually underlined for readability.

We show example system summaries in Table 6. The abstractive summaries generated by BART-Large are fluent but their content lacks specificity, so are the summaries produced by LexRank and Quantized Transformer. Particularly, QT does not seem to perform well on this task despite that the model has been retrained on livestream transcripts.<sup>8</sup> We believe this is partly because words and phrases tend to repeat themselves in review documents, and while spoken utterances are verbose, there is little repetition found in the transcripts. We observe that summary segments selected by FluCovRank are on-topic but they are ungrammatical and difficult to

<sup>8</sup><https://github.com/stangelid/qt/blob/main/custom.md>

Utterances	C1	C2	C3
0 Hello good morning everybody welcome high foster highly art.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
1 <del>Hi Lisa, welcome everyone.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
2 I hope you guys are having a good day so far.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
3 Good to see you were going to be doing cloud studies today.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
4 So if anybody is interested in joining in, if you want to work on some skies for your landscapes for future landscapes, this is what we're going to be doing.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
5 <del>Photo studies of today.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6 So I'm going to be painting from these images and these beautiful photos are from various photographers.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
7 You can find links to their original posts below.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
8 The stream in the description.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
9 One of them is from Morguefile, One is from unsplash, well, two are from Unsplash and one is from pixels there a little bit from all over the place, but you can find the photographers below if you'd like.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
10 Hey Jennifer, saw the images.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
11 I really love the top left and bottom right.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
12 The colors are beautiful but I'm partial through colorful Sky scenes.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
13 <del>Yeah, I totally agree.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
14 Let's see top left, bottom right.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
15 Those yeah well top right also is like very Contra high contrast that tends to like grab my attention when I look at the sheet but I would say top left and bottom right give me the most like happy feels.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
...			

Table 7: A snippet from *Digital Painting Studies with Maddy Bellwoar-Clouds*. We show the most prominent latent codes and their representative utterances ('X'). Human annotated summary utterances are colored gray and ultra-short utterances are crossed out.

interpret without context. In contrast, StreamHover can identify on-topic and informative utterances related to digital painting. We provide more examples in the supplementary materials.

In Table 7, we study the most prominent latent codes (C1-3) and their associated utterances. We define representative utterances as those frequently assigned to these codes (Eq. (3)). We observe that C1 usually contains a skewed number of utterances that are commonly seen in the data and not representative of the input; C2 contains lengthy but not necessarily summary-worthy utterances. In our experiments, we exclude C1/C2 before performing grid search on all codes to find the set of prominent codes: we use  $P=50$  tuned on the valid set which is effective in helping identify summary utterances.<sup>9</sup>

We conduct a human evaluation to assess how StreamHover compares to strong extractive and abstractive baselines. They are (a) LexRank (Erkan and Radev, 2004), (b) FluCovRank (Shang et al., 2018) and (c) BART-Large (Lewis et al., 2020); the latter two are abstractive systems. Each evaluator is shown a video clip with a synchronized display of the transcript followed by four system summaries, shown in random order to remove any positional

<sup>9</sup>For increased training stability of variational autoencoder (VAE) models, we refer the reader to (Li et al., 2019a).



bias. The evaluator is asked to select the best and worst of the summaries according to each of these criteria: *Fluency/Coherence*: is the the summary well-presented, grammatically correct and easy to read? *Informativeness*: does the summary provide useful information about the video clip? *Overall Quality*: is the summary of good quality considering both content and linguistic aspects?

We randomly sample 100 clips from the test set. Each clip and its summaries are judged by five evaluators that we recruit from Amazon mechanical turk.<sup>10</sup> Table 5 shows the performance of all systems measured by Best-Worst Scaling (Kiritchenko and Mohammad, 2016), where the score of a system is computed as the percentage of times it was selected as the best minus the worst. The range of scores is [-1,1]. Figure 5 shows how frequently a system is chosen to produce the “best summary.” We observe that StreamHover achieves an overall score of 0.52 and it is selected as the best summary in over half of the times.

## 6 Conclusion

We present StreamHover, a new framework for annotating and summarizing livestream transcripts. Our dataset contains over 500 hours of videos annotated with extractive and abstractive summaries. We explored an extractive method leveraging VQ-VAE to identify salient summary utterances and obtained strong results. Future work includes boosting summarization solutions to provide users a concentrated overview of streamed content.

## References

- Stefanos Angelidis, Reinald Kim Amplayo, Yoshihiko Suhara, Xiaolan Wang, and Mirella Lapata. 2021. [Extractive opinion summarization in quantized transformer spaces](#). *Transactions of the Association for Computational Linguistics (TACL)*.
- Ron Artstein and Massimo Poesio. 2008. [Survey article: Inter-coder agreement for computational linguistics](#). *Computational Linguistics*.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#).
- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. [Estimating or propagating gradients through stochastic neurons for conditional computation](#).
- Levent Besik. 2020. Google named a leader in the Gartner 2020 magic quadrant for cloud AI developer services. In <https://cloud.google.com/blog/products/ai-machine-learning/google-named-leader-gartner-2020-magic-quadrant-cloud-ai-developer-services>.
- Arthur Bražinskas, Mirella Lapata, and Ivan Titov. 2020. [Few-shot learning for opinion summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4119–4135, Online. Association for Computational Linguistics.
- Ziqiang Cao, Furu Wei, Wenjie Li, and Sujian Li. 2017. [Faithful to the original: Fact aware neural abstractive summarization](#). *CoRR*, abs/1711.04434.
- Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, Guillaume Lathoud, Mike Lincoln, Agnes Lisowska, Iain McCowan, Wilfried Post, Dennis Reidsma, and Pierre Wellner. 2006. The ami meeting corpus: A pre-announcement. In *Machine Learning for Multimodal Interaction*, pages 28–39, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Mingda Chen, Zewei Chu, Sam Wiseman, and Kevin Gimpel. 2021. [Summscreen: A dataset for abstractive screenplay summarization](#). *CoRR*, abs/2104.07091.
- Yen-Chun Chen and Mohit Bansal. 2018. [Fast abstractive summarization with reinforce-selected sentence rewriting](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–686, Melbourne, Australia. Association for Computational Linguistics.
- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. [A discourse-aware attention model for abstractive summarization of long documents](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 615–621, New Orleans, Louisiana. Association for Computational Linguistics.
- Yue Dong, Yikang Shen, Eric Crawford, Herke van Hoof, and Jackie Chi Kit Cheung. 2018. [Bandit-Sum: Extractive summarization as a contextual bandit](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3739–3748, Brussels, Belgium. Association for Computational Linguistics.
- Günes Erkan and Dragomir R. Radev. 2004. [LexRank: Graph-based lexical centrality as salience in text summarization](#). *Journal of Artificial Intelligence Research*.

<sup>10</sup>They are required to be Master workers, have an approval rate >90% and completed >1K tasks, and located in the U.S.

- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. [Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics.
- Sebastian Gehrmann, Yuntian Deng, and Alexander Rush. 2018. [Bottom-up abstractive summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4098–4109, Brussels, Belgium. Association for Computational Linguistics.
- J. J. Godfrey, E. C. Holliman, and J. McDaniel. 1992. [Switchboard: telephone speech corpus for research and development](#). In *[Proceedings] ICASSP-92: 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 517–520 vol.1.
- Matt Grenander, Yue Dong, Jackie Chi Kit Cheung, and Annie Louis. 2019. [Countering the effects of lead bias in news summarization via multi-stage training and auxiliary losses](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6019–6024, Hong Kong, China. Association for Computational Linguistics.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Estève. 2018. [Ted-lium 3: Twice as much data and corpus repartition for experiments on speaker adaptation](#). *Lecture Notes in Computer Science*, page 198–208.
- A. Janin, D. Baron, J. Edwards, D. Ellis, D. Gelbart, N. Morgan, B. Peskin, T. Pfau, E. Shriberg, A. Stolcke, and C. Wooters. 2003. The icsi meeting corpus. In *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03)*, volume 1, pages I–I.
- Shuning Jin, Sam Wiseman, Karl Stratos, and Karen Livescu. 2020. [Discrete latent variable representations for low-resource text classification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4831–4842, Online. Association for Computational Linguistics.
- Sushant Kafle and Matt Huenerfauth. 2018. [A corpus for modeling word importance in spoken dialogue transcripts](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Chris Kedzie, Kathleen McKeown, and Hal Daumé III. 2018. [Content selection in deep learning models of summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1818–1828, Brussels, Belgium. Association for Computational Linguistics.
- Svetlana Kiritchenko and Saif M. Mohammad. 2016. [Capturing reliable fine-grained sentiment associations by crowdsourcing and best–worst scaling](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 811–817, San Diego, California. Association for Computational Linguistics.
- Jia Jin Koay, Alexander Roustai, Xiaojin Dai, Dillon Burns, Alec Kerrigan, and Fei Liu. 2020. [How domain terminology affects meeting summarization performance](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5689–5695, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jia Jin Koay, Alexander Roustai, Xiaojin Dai, and Fei Liu. 2021. [A sliding-window approach to automatic creation of meeting minutes](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 68–75, Online.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 540–551, Hong Kong, China. Association for Computational Linguistics.
- Faisal Ladhak, Bryan Li, Yaser Al-Onaizan, and Kathleen McKeown. 2020. [Exploring content selection in summarization of novel chapters](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5043–5054, Online. Association for Computational Linguistics.
- Logan Lebanoff, John Muchovej, Franck Dernoncourt, Doo Soon Kim, Seokhwan Kim, Walter Chang, and Fei Liu. 2019. [Analyzing sentence fusion in abstractive summarization](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 104–110, Hong Kong, China. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Bohan Li, Junxian He, Graham Neubig, Taylor Berg-Kirkpatrick, and Yiming Yang. 2019a. [A surprisingly effective fix for deep latent variable modeling of text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*

- and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 3603–3614, Hong Kong, China. Association for Computational Linguistics.
- Manling Li, Lingyu Zhang, Heng Ji, and Richard J. Radke. 2019b. [Keep meeting summaries on topic: Abstractive multi-modal meeting summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2190–2196, Florence, Italy. Association for Computational Linguistics.
- Mingzhe Li, Xiuying Chen, Shen Gao, Zhangming Chan, Dongyan Zhao, and Rui Yan. 2020. [VMSMO: Learning to generate multimodal summary for video-based news articles](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9360–9369, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Feifan Liu and Yang Liu. 2008. [Correlation between ROUGE and human evaluation of extractive meeting summaries](#). In *Proceedings of ACL-08: HLT, Short Papers*, pages 201–204, Columbus, Ohio. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2019. [Text summarization with pretrained encoders](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Matthew Marge, Satantjeev Banerjee, and Alexander Rudnicky. 2010. [Using the Amazon Mechanical Turk to transcribe and annotate meeting speech for extractive summarization](#). In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 99–107, Los Angeles. Association for Computational Linguistics.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Rada Mihalcea and Paul Tarau. 2004. [TextRank: Bringing order into text](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain. Association for Computational Linguistics.
- Gabriel Murray and Giuseppe Carenini. 2008. [Summarizing spoken and written conversations](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 773–782, Honolulu, Hawaii. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018a. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018b. [Ranking sentences for extractive summarization with reinforcement learning](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1747–1759, New Orleans, Louisiana. Association for Computational Linguistics.
- Shruti Palaskar, Jindřich Libovický, Spandana Gella, and Florian Metze. 2019. [Multimodal abstractive summarization for how2 videos](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6587–6596, Florence, Italy. Association for Computational Linguistics.
- Pinelopi Papalampidi, Frank Keller, Lea Frermann, and Mirella Lapata. 2020. [Screenplay summarization using latent narrative structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1920–1933, Online. Association for Computational Linguistics.
- Gerald Penn and Xiaodan Zhu. 2008. [A critical re-assessment of evaluation baselines for speech summarization](#). In *Proceedings of ACL-08: HLT*, pages 470–478, Columbus, Ohio. Association for Computational Linguistics.
- Gabriele Prato, Ella Charlaix, and Mehdi Rezagholizadeh. 2020. [Fully quantized transformer for machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1–14, Online. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text](#)

- generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Guokan Shang, Wensi Ding, Zekun Zhang, Antoine Tixier, Polykarpos Meladianos, Michalis Vazirgiannis, and Jean-Pierre Lorré. 2018. **Unsupervised abstractive meeting summarization with multi-sentence compression and budgeted submodular maximization**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 664–674, Melbourne, Australia. Association for Computational Linguistics.
- Kaiqiang Song, Bingqing Wang, Zhe Feng, and Fei Liu. 2021. **A new approach to overgenerating and scoring abstractive summaries**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1392–1404, Online.
- Jiwei Tan, Xiaojun Wan, and Jianguo Xiao. 2017. **Abstractive document summarization with a graph-based attentional neural model**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1171–1181, Vancouver, Canada. Association for Computational Linguistics.
- Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. 2017. **Neural discrete representation learning**. In *Advances in Neural Information Processing Systems*, volume 30, pages 6306–6315. Curran Associates, Inc.
- Lucy Vanderwende, Hisami Suzuki, Chris Brockett, and Ani Nenkova. 2007. **Beyond SumBasic: Task-focused summarization with sentence simplification and lexical expansion**. *Information Processing and Management*, 43(6):1606–1618.
- Danqing Wang, Pengfei Liu, Yining Zheng, Xipeng Qiu, and Xuanjing Huang. 2020. **Heterogeneous graph neural networks for extractive document summarization**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6209–6219, Online. Association for Computational Linguistics.
- Lu Wang and Claire Cardie. 2013. **Domain-independent abstract generation for focused meeting summarization**. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 1395–1405, Sofia, Bulgaria.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Wen Xiao and Giuseppe Carenini. 2019. **Extractive summarization of long documents by combining global and local context**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3011–3021, Hong Kong, China. Association for Computational Linguistics.
- Jiacheng Xu and Greg Durrett. 2019. **Neural extractive text summarization with syntactic compression**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3292–3303, Hong Kong, China. Association for Computational Linguistics.
- Michihiro Yasunaga, Rui Zhang, Kshitijh Meelu, Ayush Pareek, Krishnan Srinivasan, and Dragomir Radev. 2017. **Graph-based neural multi-document summarization**. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 452–462, Vancouver, Canada. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. **Bertscore: Evaluating text generation with bert**.
- Chen Zhao, Chenyan Xiong, Corby Rosset, Xia Song, Paul Bennett, and Saurabh Tiwary. 2020. **Transformer-xh: Multi-evidence reasoning with extra hop attention**. In *International Conference on Learning Representations*.
- Ming Zhong, Da Yin, Tao Yu, Ahmad Zaidi, Mutethia Mutuma, Rahul Jha, Ahmed Hassan Awadallah, Asli Celikyilmaz, Yang Liu, Xipeng Qiu, and Dragomir Radev. 2021. **QMSum: A new benchmark for query-based multi-domain meeting summarization**. In *Proceedings of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Chenguang Zhu, Yang Liu, Jie Mei, and Michael Zeng. 2021. **MediaSum: A large-scale media interview dataset for dialogue summarization**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5927–5934, Online.
- Junnan Zhu, Haoran Li, Tianshang Liu, Yu Zhou, Jiajun Zhang, and Chengqing Zong. 2018. **MSMO: Multimodal summarization with multimodal output**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4154–4164, Brussels, Belgium. Association for Computational Linguistics.

## A Bēhance Dataset

We collect a total of 5,398 streamed videos from Behance.net. Some streamers opt-out of the transcription service provided by Microsoft Automatic Speech Recognition, so transcripts are not available for these videos. We create a list of domain keywords by finding 50 most frequently appearing words from video titles (stopwords are excluded). Examples include ‘fresco’, ‘adobe’, ‘photoshop’, ‘illustration’, ‘art’, ‘painting’, ‘drawing’, ‘illustrator’, ‘character’, ‘design.’ The keywords are used to select videos for human annotation. 2,360 videos have transcripts available and contain at least one of our domain keywords in their titles. These videos are split into clips of 5-minute each. Some clips contain little or no verbal content. We thus remove clips that contain very few words ( $\leq 333$  words) or utterances ( $\leq 38$  utterances). These thresholds are determined using the average values of all clips. Videos with less than 5 valid clips are also removed from consideration. This preprocessing step gives 6,003 clips from 381 videos. During annotation, our annotators find 582 clips to contain only chit-chats, suggesting that these clips are uninformative. 11 videos contain only chit-chat clips, they are subsequently removed from the dataset, yielding a total of 5,421 clips from 370 videos that are split into train, validation and test sets.

## B Baseline Summarizers

Our neural abstractive baselines include pre-trained BART-large (Lewis et al., 2020), BART-large-cnn, and T5-large (Raffel et al., 2020). We follow the HuggingFace implementation (Wolf et al., 2020). Utterances that are longer than 5 words are concatenated into a flat sequence, which is used as the input to each summarizer. The model parameters include: the maximum and minimum summary lengths are 150 and 15 tokens, respectively. We use a beam size of 5 with early stopping. The length penalty is 1.0. “no\_repeat\_ngram\_size” is set to 3, such that a trigram cannot occur more than once in the summary.

Our extractive baselines include LexRank (Erkan and Radev, 2004), TextRank (Mihalcea and Tarau, 2004), and SumBasic (Vanderwende et al., 2007). They are implemented using the Sumy library where we adopt the default text parser and stemmer. Our unsupervised summarizer for speech transcript summarization (Shang et al., 2018) uses the following settings: we report the FluCovRank

scores. The number of components used in LSA is 25. The number of utterance communities is 35. The number of clusters is 6, with a scaling factor of 1.3 and lambda of 0.4. The size of the summary is set to 50 words.

## C Example Summaries

We show example summaries generated by different summarizers: FluConvRank (Shang et al., 2018), LexRank (Erkan and Radev, 2004), BART-large (Lewis et al., 2020) and StreamHover. We also show the top-3 most prominent latent codes and their associated utterances. We choose 5 representative utterances for each code that are most frequently assigned to this code. We observe that C1 utterances are frequently seen in the data (chit-chats) and not representative of the input. C2 is associated with lengthy but not necessarily summary-worthy utterances. C3 utterances are both comprehensive and contain diverse information. In our experiments, we exclude C1/C2 before performing grid search on all codes to find the set of prominent codes. It allows us to effectively identify summary utterances without biasing towards the lengthy ones.

FluCovRank
<ul style="list-style-type: none"> <li>digital painting today / 3 hours minimum / explain my painting process / starbucks florida cup / really really cheerful scene / prob hey welcome hi kumar / canvas sizes 3000 by 2000 pixels / sore throat so todays stream might be a tiny bit shorter / otherwise just sit back and relax and enjoy the painting / find</li> </ul>
LexRank
<ul style="list-style-type: none"> <li>If you guys are interested in painting along with me or doing a little sketch or whatever, you have time for your always welcome to my reference.</li> <li>Is what I have planned to work on this if I'm feeling good, we might go on or do another painting, but I'm drinking some tea with honey and I have to kind of be careful with my voice a little bit talking so much, but I'm otherwise I'm good.</li> <li>So again, that's where you can find my references below in the stream description.</li> <li>You guys are welcome to paint along if you're interested in doing that.</li> <li>I've lived in a few different states, but where I spent most of my life was in Florida so.</li> </ul>
BART-Large
<ul style="list-style-type: none"> <li>Good afternoon, good evening, good morning, good day, good afternoon, how are you? Good morning, hi Mohammed, hi kumar, hi muhammad, hi hi kumar, hi hello hi hi hi Mohammed. I hope you've had a good day so far. I'm doing pretty good pro to be honest. I have a little bit of a sore throat, so today's stream might be a tiny bit shorter than some of the others. It's still going to be around 3 hours minimum. Is what I have planned to work on this if I'm feeling good, we might go on or do another painting, but I'm drinking some tea with honey and I have to kind of</li> </ul>
Quantized Transformer
<ul style="list-style-type: none"> <li>We're going to be doing some digital painting today.</li> <li>As usual, I'm going to be working from this photo as a reference.</li> <li>I hope you're all doing well.</li> <li>Image is below in the stream description and above the chat.</li> <li>It's still going to be around 3 hours minimum.</li> <li>Is what I have planned to work on this if I'm feeling good, we might go on or do another painting, but I'm drinking some tea with honey and I have to kind of be careful with my voice a little bit talking so much, but I'm otherwise I'm good.</li> <li>Um prob hey welcome Hi Kumar.</li> <li>You guys are welcome to paint along if you're interested in doing that.</li> <li>I hope it's going to be good.</li> <li>My canvas sizes 3000 by 2000 pixels.</li> </ul>
StreamHover (Ours)
<ul style="list-style-type: none"> <li>It's like a little cafe table in Spain, with some dappled lighting coming through these trees.</li> <li>I really like the pattern of light on the building over here, and the just overall feeling the different colors.</li> <li>If you guys are interested in painting along with me or doing a little sketch or whatever, you have time for your always welcome to my reference.</li> <li>Is what I have planned to work on this if I'm feeling good, we might go on or do another painting, but I'm drinking some tea with honey and I have to kind of be careful with my voice a little bit talking so much, but I'm otherwise I'm good.</li> <li>I am using pure ref to put the little reference image up here so we can see a thumbnail view of it while we're painting.</li> </ul>

Table 8: Example system summaries from *Digital Painting with Maddy Bellwoar*. BART summary is fluent but its content lacks specificity, as is the case for LexRank. Summary segments selected by FluCovRank are ungrammatical. StreamHover identifies on-topic and informative utterances related to digital painting.

Utterances	C1	C2	C3
0 <del>Hi everybody, good morning.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
1 Good afternoon, good evening whatever time it is for you.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
2 <del>Thanks for being here.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
3 <del>Hi Fernando Allison, Welcome.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
4 I hope you've had a good day so far.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
5 I hope you've all been well.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6 If I haven't seen you since last week, I hope you had a good weekend.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
7 We're going to be doing some digital painting today.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
8 As usual, I'm going to be working from this photo as a reference.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
9 I think it's really beautiful.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
10 It's like a little cafe table in Spain, with some dappled lighting coming through these trees.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
11 I really like the pattern of light on the building over here, and the just overall feeling the different colors.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
12 the Blues and pinks and yellows on the wall here.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
13 It's really, really cheerful scene.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
14 Hi Mohammed, high pro, how are you?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
15 I hope you're all doing well.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
16 If you guys are interested in painting along with me or doing a little sketch or whatever, you have time for your always welcome to my reference.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
17 Image is below in the stream description and above the chat.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
18 Also, if you check the info panel it should be there as well so this photo is from Unsplash.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
19 You can find a link to the original post from the photographer below.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
20 <del>The stream.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
21 I'm doing pretty good pro to be honest.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
22 I have a little bit of a sore throat, so today's stream might be a tiny bit shorter than some of the others.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
23 It's still going to be around 3 hours minimum.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
24 Is what I have planned to work on this if I'm feeling good, we might go on or do another painting, but I'm drinking some tea with honey and I have to kind of be careful with my voice a little bit talking so much, but I'm otherwise I'm good.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
25 Um prob hey welcome Hi Kumar.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
26 So again, that's where you can find my references below in the stream description.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
27 You guys are welcome to paint along if you're interested in doing that.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
28 Otherwise, just sit back and relax and enjoy the painting.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
29 I'll try to do a nice job with this one.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
30 I hope it's going to be good.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
31 I really like the reference so.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
32 <del>Hi Dante.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
33 So I'm going to explain my painting process as we go through here.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
34 And if you guys have any questions you can let me know.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
35 I am using pure ref to put the little reference image up here so we can see a thumbnail view of it while we're painting.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
36 My canvas sizes 3000 by 2000 pixels.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
37 <del>If you're curious.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
38 Yes, I'm using a different Cup today.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
39 This is my Starbucks Florida Cup.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
40 Maybe more information than you need, but yeah, so I'm I grew up in Florida.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
41 I've lived in a few different states, but where I spent most of my life was in Florida so.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
42 I don't live there now.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Table 9: A transcript snippet from *Digital Painting with Maddy Bellwoar*. We show the most prominent latent codes and their representative utterances ('X'). Human annotated summary utterances are colored gray and ultra-short utterances are crossed out.

FluCovRank
<ul style="list-style-type: none"> <li>view is created by them stitching images / maps years was a fitness studio yeah / image stitching just has some weird that happens team / oh if you want to share a link for map crunch / leave so weird parts / hope achieve that feeling of the water is really moving / yeah its</li> </ul>
LexRank
<ul style="list-style-type: none"> <li>So I'm going to start to work on the foreground now and this is going to be tricky, so this is one of the things that made me choose to paint this reference.</li> <li>The rushing water in the foreground kind of coming towards us and the ripples that it creates an the little bubbles and all that kind of stuff.</li> <li>Sometimes the image stitching just has some weird stuff that happens, Oh team.</li> <li>Oh, if you want to share a link for map crunch.</li> <li>Alright that should work this time.</li> </ul>
BART-Large
<ul style="list-style-type: none"> <li>Trying to figure out those relationships near the end or like when you're trying to like if you're going to try to get the one or like if the one is like if I'm working on the one I'm going to do it's like if it's not working. I'm ready for some dog like with half his body has been overlapped by a boat. Yeah, it's always really satisfying the color combination in these kind of desert scenes, the nice blue Sky and like the Warm Reds and oranges. It's like you get the extra bonus of like some purple thrown in there and that's just perfect. You know, I've seen people like riding camels and stuff, he's taking pictures</li> </ul>
Quantized Transformer
<ul style="list-style-type: none"> <li>So I hope I can achieve that feeling of the water is really moving there.</li> <li>You have to click the share button and then copy the link that is there if you copy the link from the URL above it's not going to send the right image.</li> <li>So what you can do is click right here where it says share and then copy this.</li> <li>OK, so these so right now, I'm just going to start with a darker shadows.</li> <li>Of where I see the ripples.</li> <li>Alright that should work this time.</li> <li>You know where some of the really, really fantastic ones come from.</li> <li>But people will be on the boat on foot.</li> <li>Maps years was a fitness studio, yeah, there's all kinds of stuff.</li> <li>I've painted with references with stuff like that.</li> </ul>
StreamHover (Ours)
<ul style="list-style-type: none"> <li>So I'm going to start to work on the foreground now and this is going to be tricky, so this is one of the things that made me choose to paint this reference.</li> <li>The rushing water in the foreground kind of coming towards us and the ripples that it creates an the little bubbles and all that kind of stuff.</li> <li>Yeah, it's always really satisfying the color combination in these kind of desert scenes, the nice blue Sky and like the Warm Reds and oranges.</li> <li>I'm ready for some dog like with half his body has been overlapped by a boat.</li> <li>People are in a car traveling there, you get a lot of side of the road like kind of bland images, but then you also get like the most amazing pictures and people also upload images and stuff, I think that might be.</li> </ul>

Table 10: Example system summaries from *Virtual Plein Air Landscape Painting*. The BART summary is fluent but its content lacks specificity, as is the case for LexRank. The summary segments selected by FluCovRank are ungrammatical. StreamHover identifies on-topic and informative summary utterances.

Utterances	C1	C2	C3
0 Trying to figure out those relationships near the end or like force them if they're not working.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
1 So I'm going to start to work on the foreground now and this is going to be tricky, so this is one of the things that made me choose to paint this reference.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
2 'cause I wanted to figure out how to paint this kind of situation.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
3 The rushing water in the foreground kind of coming towards us and the ripples that it creates an the little bubbles and all that kind of stuff.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
4 So I hope I can achieve that feeling of the water is really moving there.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
5 Might give some inspiration alright let's see got an artstation link.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6 <del>Oh, cool.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
7 <del>Good vibes good vibes.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
8 <del>Other desert desert scene.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
9 Yeah, it's always really satisfying the color combination in these kind of desert scenes, the nice blue Sky and like the Warm Reds and oranges.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
10 I like the one I'm working on 2 'cause.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
11 It's like you get the extra bonus of like some purple thrown in there and that's just perfect.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
12 <del>For me.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
13 <del>Yeah, that's cool thanks.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
14 Team I says, I was looking for a nice location on map French and I found this alright.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
15 I'm ready for some dog like with half his body has been overlapped by a boat.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
16 Sometimes the image stitching just has some weird stuff that happens, Oh team.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
17 Oh, if you want to share a link for map crunch.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
18 You have to click the share button and then copy the link that is there if you copy the link from the URL above it's not going to send the right image.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
19 So what you can do is click right here where it says share and then copy this.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
20 And then that will be the right one, 'cause I think it sent the wrong one from you.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
21 <del>Yeah.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
22 <del>Alright cool.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
23 Will get back to painting?	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
24 OK, so these so right now, I'm just going to start with a darker shadows.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
25 Of where I see the ripples.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
26 I'm not worried about putting like every ripple in the right spot.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
27 But we couldn't get the overall feel this.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
28 <del>Some more Brown here.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
29 Alright that should work this time.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
30 What's wrong with that picture?	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
31 <del>Does actually seems nice?</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
32 I don't see anything wrong with that landscape Timo?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
33 When you when you're using map crunch you're going to get a lot of like side of the road road views because you know if you think about where the car is going.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
34 Order fact that it's your in the Google Maps.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
35 People are in a car traveling there, you get a lot of side of the road like kind of bland images, but then you also get like the most amazing pictures and people also upload images and stuff, I think that might be.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
36 You know where some of the really, really fantastic ones come from.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
37 But people will be on the boat on foot.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
38 You know, I've seen people like riding camels and stuff, he's taking pictures for Google.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
39 Maps years was a fitness studio, yeah, there's all kinds of stuff.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
40 Yeah, but sometimes you get something weird because the 3D or the 360 Panorama.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
41 View is created by them stitching images together.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
42 So, sometimes you have weird situations where stuff overlaps.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Table 11: A snippet from *Virtual Plein Air Landscape Painting*. We show the most prominent latent codes and their representative utterances ('X'). Human annotated summary utterances are colored gray and ultra-short utterances are crossed out.

FluCovRank
<ul style="list-style-type: none"> <li>drawing diesel letters / surface pen 4 i have the 3 / latest song version of the surface / pin so i got a surface pen / uhm this pan to check / avoid these artifacts in astana / file ready an complete im just listening / early in october 20th 19th / bump on phone so yeah</li> </ul>
LexRank
<ul style="list-style-type: none"> <li>Thank you for sticking around wanna show you some real cool stuff, not right there.</li> <li>Bump on Phone so yeah, that's the wrong letter, but let's go.</li> <li>But I thought we fresco does not have good, um Palm rejection and so when you using the pen and drawing whether you're on the iPad or surface.</li> <li>We go come on come here pin so I got a Surface Pen here.</li> <li>Pan's had details that were missing right now, I'm using the vector pen the vector brushes are in.</li> </ul>
BART-Large
<ul style="list-style-type: none"> <li>Thank you for sticking around wanna show you some real cool stuff, not right there. I'm on the Surface Pro. That's the latest song version of the surface. It just came out early in October 20th 19th. The app and it's working it's a work in progress OK is constantly being updated with the features and bugs. They sent me uhm this pan to check out. Pan's had details that were missing right now, I'm using the vector pen the vector brushes are in. So I'm doing this little dots and what happens is we on with Adobe fresco is there. But I thought we fresco does not have good, um Palm rejection and so when you using the</li> </ul>
Quantized Transformer
<ul style="list-style-type: none"> <li>I do have my Surface Pen somewhere it's up there.</li> <li>I'm on the Surface Pro.</li> <li>That's the latest song version of the surface.</li> <li>You end up with all these little artifacts that when your Palm lays on the screen.</li> <li>And so as I was drawing diesel letters.</li> <li>I got to the point where.</li> <li>I just cannot avoid these artifacts in Astana kept them.</li> <li>I need to do some of these lines here is geared is closeout all.</li> <li>If this doesn't have a pain.</li> <li>Just did good thing is I have extra pants, not only do I have all here's one, not only do I have extra pants.</li> <li>OK, here, we go ever it is yeah, some of these.</li> </ul>
StreamHover (Ours)
<ul style="list-style-type: none"> <li>It's gotta stay and the same letter so they was the letter C and I'm just using my hands my fingers.</li> <li>The app and it's working it's a work in progress OK is constantly being updated with the features and bugs.</li> <li>But I thought we fresco does not have good, um Palm rejection and so when you using the pen and drawing whether you're on the iPad or surface.</li> <li>I kept them in the drawing and now they are part of the look of each of these letters, which is crazy.</li> <li>Just did good thing is I have extra pants, not only do I have all here's one, not only do I have extra pants.</li> </ul>

Table 12: Example system summaries from *Creating ABC Childrens Book Art on Adobe Fresco and Adobe Illustrator Part 18*. BART summary is fluent but its content lacks specificity, as is the case for LexRank. Summary segments selected by FluCovRank are ungrammatical. StreamHover identifies on-topic and informative summary utterances.

Utterances	C1	C2	C3
0 Thank you for sticking around wanna show you some real cool stuff, not right there.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
1 It's gotta stay and the same letter so they was the letter C and I'm just using my hands my fingers.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
2 I do have my Surface Pen somewhere it's up there.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
3 I'm on the Surface Pro.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
4 <del>7</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
5 That's the latest song version of the surface.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6 It just came out early in October 20th 19th.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
7 I know that in a few days were gonna be in 2020, which is why I want to have.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
8 This file ready an complete I'm just listening.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
9 Bump on Phone so yeah, that's the wrong letter, but let's go.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
10 <del>It's just be gone.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
11 So I'm doing this little dots and what happens is we on with Adobe fresco is there.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
12 It doesn't do a good job.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
13 The app and it's working it's a work in progress OK is constantly being updated with the features and bugs.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
14 But I thought we fresco does not have good, um Palm rejection and so when you using the pen and drawing whether you're on the iPad or surface.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
15 You end up with all these little artifacts that when your Palm lays on the screen.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
16 And so as I was drawing diesel letters.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
17 I got to the point where.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
18 I just cannot avoid these artifacts in Astana kept them.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
19 I kept them in the drawing and now they are part of the look of each of these letters, which is crazy.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
20 And then there's some of these letters that don't have detail like this one.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
21 So now I am going to use my pen there.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
22 We go come on come here pin so I got a Surface Pen here.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
23 <del>Anthony my glasses.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
24 But the lightless terrible OK here we go let me.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
25 I need to do some of these lines here is geared is closeout all.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
26 <del>Define.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
27 If this doesn't have a pain.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
28 Come on, it does have a pen.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
29 <del>Usually.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
30 <del>I guess the battery.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
31 Just did good thing is I have extra pants, not only do I have all here's one, not only do I have extra pants.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
32 <del>There we go.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
33 I am I have the Surface Pen 4 I have the 3.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
34 <del>I have.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
35 I also have other pins with other brands like the Wacom, the bamboo adanet and this other.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
36 <del>Pan brand.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
37 They sent me uhm this pan to check out.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
38 <del>They mean the company.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
39 At the Renaissance so it's pretty cool.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
40 OK, here, we go ever it is yeah, some of these.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
41 Pan's had details that were missing right now, I'm using the vector pen the vector brushes are in.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
42 <del>There we go.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
43 I'ma try not to yell but I know I can hear.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
44 Just a big bounce in my voice and Missy Volume.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
45 OK, so now we're back we're back 'cause I could hear the sound an it was very loud.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
46 I'm sorry I already have a loud voice.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Table 13: A transcript snippet from *Creating ABC Childrens Book Art on Adobe Fresco and Adobe Illustrator Part 18*. We show the most prominent latent codes and their representative utterances ('X'). Human annotated summary utterances are colored gray and ultra-short utterances are crossed out.



<p><b>FluCovRank</b></p> <ul style="list-style-type: none"> <li>learning process proceed / draw this out and post it where people / sit down an just trying to draw a picture / spent like 3 hours / alot like ill draw something like old man i want to draw / consistently draw this character and erase and restart and you cant necessarily / practice will</li> </ul>
<p><b>LexRank</b></p> <ul style="list-style-type: none"> <li>The practice will do a lot more for you than drawing a picture.</li> <li>You might get one picture that might turn out really cool and you can post it and stuff like that, but if you work on your practice and you get that practice the way that you want it, you can use that practice forever for more than one picture.</li> <li>You stop trying to make the picture and get back to practice.</li> <li>So if there's things that you need to practice and stuff like that, you stand to gain a lot more from that practice then.</li> <li>Sit down an just trying to draw a picture.</li> </ul>
<p><b>BART-Large</b></p> <ul style="list-style-type: none"> <li>I do I like to do more practice in anything. Practices gives me this feel that I've never had when I was younger. You know, of course, everybody will tell you. You might get one picture that might turn out really cool and you can post it and stuff like that, but if you work on your practice and you get that practice the way that you want it, you can use that practice forever for more than one picture. So that's generally would always tell myself when I'm drawing, and I'm like. Oh crap, here I go again. I should draw this into a picture and then I end up skipping practice to save whatever I'm doing as a picture so that I can. Just 'cause</li> </ul>
<p><b>Quantized Transformer</b></p> <ul style="list-style-type: none"> <li>Trying to make a picture, right?</li> <li>And then I'm like OK alright I gotta I gotta center myself You know I gotta always gotta check myself every now and again or back that was kind of how it wasn't back then but now I don't really like to draw pictures.</li> <li>I do I like to do more practice in anything.</li> <li>You sit up and you put on like references for the superhero be in movies or something like that.</li> <li>You pull up references on whatever social media or wherever you get your images from.</li> <li>Whereas if you were just trying to practice, you would realize that all of that stuff that you're about to throw away or disregard was actually what you were actually searching for.</li> </ul>
<p><b>StreamHover (Ours)</b></p> <ul style="list-style-type: none"> <li>Doesn't matter what styles you draw, so don't overly complicated by like Oh my God, you know now you could draw it.</li> <li>I used to do that a lot when I was younger where I would do something and be like holy crap, didn't think that would come out that way.</li> <li>And then I'm like OK alright I gotta I gotta center myself You know I gotta always gotta check myself every now and again or back that was kind of how it wasn't back then but now I don't really like to draw pictures.</li> <li>They don't see it happening until it's too late where it's like, Oh Man, I could totally practice this right now, or I could Draw Something Super Duper cool, you know, a piece of fan art or something like that.</li> <li>You're going to struggle with for a couple of days or couple minutes and then figure out you can't draw it or you have a hard time drawing it and then just giving up.</li> </ul>

Table 14: Example system summaries from *Call me Derek: Value Study*. The BART summary is fluent but its content lacks specificity, as is the case for LexRank. Summary segments selected by FluCovRank are ungrammatical. StreamHover identifies on-topic and informative summary utterances.

Utterances	C1	C2	C3
0 And then once I get my general OK this is what the hairstyle is gonna be right here.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
1 I definitely want to condom mark the flow.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
2 Well, I'll do that later.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
3 And now I just kind of relax and just.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
4 Draw hairstyles as long as there as long as his long hair.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
5 Doesn't matter what styles you draw, so don't overly complicated by like Oh my God, you know now you could draw it.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6 <del>Now you can draw.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
7 <del>Faces and stuff.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
8 If you want, I try not to draw too much of the face.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
9 Just 'cause I don't want to.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
10 I don't want to get caught up in that thing that I always do where when I'm drawing a face or I'm doing practice.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
11 I always do that thing where it's like holy crap you know this is a really good face.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
12 I should draw this into a picture and then I end up skipping practice to save whatever I'm doing as a picture so that I can.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
13 <del>You know?</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
14 Posted on Instagram or something like that.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
15 I used to do that a lot when I was younger where I would do something and be like holy crap, didn't think that would come out that way.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
16 I don't want to waste this in practice.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
17-21 ...	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
22 You might get one picture that might turn out really cool and you can post it and stuff like that, but if you work on your practice and you get that practice the way that you want it, you can use that practice forever for more than one picture.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
23 So that's generally would always tell myself when I'm drawing, and I'm like.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
24 Oh crap, here I go again.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
25 Trying to make a picture, right?	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
26 You stop trying to make the picture and get back to practice.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
27 And then I'm like OK alright I gotta I gotta center myself You know I gotta always gotta check myself every now and again or back that was kind of how it wasn't back then but now I don't really like to draw pictures.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
28 I do I like to do more practice in anything.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
29 Practices gives me this feel that I've never had when I was younger.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
30 When I was younger, which is so eager to get people to see my work that I wasn't necessarily thinking about the overall count, their overall.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
31 Affect that, we're going to have on.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
32 <del>My learning process proceed.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
33 And it had a big big, big big like effect on it took me longer to learn stuff because when you sit down you feel like you don't want to do anything unless you're actually drawing something to show off, which is something you don't necessarily want to get caught up in, which happens a lot of people they don't.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
35-36 ...	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
37 Get way more out of practice because practice is more focused on getting you better at art where doing a picture or whatever is cool, but it's more focused on showing off what you can do or what you already know.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
38 So if there's things that you need to practice and stuff like that, you stand to gain a lot more from that practice then.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
39 <del>You know?</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
40 Sit down an just trying to draw a picture.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
41 You're going to struggle with for a couple of days or couple minutes and then figure out you can't draw it or you have a hard time drawing it and then just giving up.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
42 Which is something that would happen alot like I'll Draw Something like old man, you know I want to draw this.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
43 Superhero, I've had a problem with drawing the superhero forever, but today is going to be the day I'm going to nail it.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Table 15: A transcript snippet from *Call me Derek: Value Study*. We show the most prominent latent codes and their representative utterances ('X'). Human annotated summary utterances are colored gray and ultra-short utterances are crossed out.

FluCovRank
<ul style="list-style-type: none"> <li>• people start off always drawing characters / favorite characters from sailor moon / mohammed roller show / yeah people are different or personalities / example prefer environment designed the character design and even though i like looking at character art by prefer to make environments / advice with a grain of salt / people specialized the</li> </ul>
LexRank
<ul style="list-style-type: none"> <li>• Some people like to have things mixed up all the time where they get bored and other people find that one thing they like and never tire of it.</li> <li>• Or what I meant is that some people start off always drawing characters, for example like they just from.</li> <li>• An anyway, some artists just like for example, love drawing characters and always draw them and never get into environments, and they go on and become a good working concepts are it could work in visual development as a character artist who knows what.</li> <li>• Just because someone specializes, it doesn't necessarily mean that they like tried all the other things first.</li> <li>• Some people sort of specialized the whole time just because that was their preference, you know.</li> </ul>
BART-Large
<ul style="list-style-type: none"> <li>• I really works good to see you. Yeah, people are different or personalities. What might work for them isn't necessarily right right for you for any number of reasons, like what David said just from person to person, those preferences can be really different. I for example, prefer environment designed the character design, and even though I like looking at character art by prefer to make environments, and that's not. I think that's why you have to take advice with a grain of salt like. See I do think people do though Omar, like maybe maybe you're not relating to it 'cause it's just not what your preferences are but. Definitely, I think it's important to seek out advice from professional artists and from people</li> </ul>
Quantized Transformer
<ul style="list-style-type: none"> <li>• It's about what you like doing.</li> <li>• I really works good to see you.</li> <li>• A specialist most of the time is someone with more experience.</li> <li>• I mean I like for example, I started drawing in like when I was a kid and when I was in middle school and high school was drawn.</li> <li>• Just because someone specializes, it doesn't necessarily mean that they like tried all the other things first.</li> <li>• This was the first study.</li> <li>• We're doing cloud studies for part of the art club were doing clouds in different scenarios, this time different.</li> <li>• This was the second one.</li> <li>• See I do think people do though Omar, like maybe maybe you're not relating to it 'cause it's just not what your preferences are but.</li> <li>• All thanks claires you thank you.</li> <li>• So I do hope that it helps.</li> </ul>
StreamHover (Ours)
<ul style="list-style-type: none"> <li>• Some people like to have things mixed up all the time where they get bored and other people find that one thing they like and never tire of it.</li> <li>• An anyway, some artists just like for example, love drawing characters and always draw them and never get into environments, and they go on and become a good working concepts are it could work in visual development as a character artist who knows what.</li> <li>• We're doing cloud studies for part of the art club were doing clouds in different scenarios, this time different.</li> <li>• See I do think people do though Omar, like maybe maybe you're not relating to it 'cause it's just not what your preferences are but.</li> <li>• I for example, prefer environment designed the character design, and even though I like looking at character art by prefer to make environments, and that's not.</li> </ul>

Table 16: Example system summaries for *Digital Painting Studies with Maddy Bellwoar-Clouds*. The BART summary is fluent but its content lacks specificity, as is the case for LexRank. The summary segments selected by FluCovRank are ungrammatical. StreamHover identifies on-topic and informative summary utterances.

Utterances	C1	C2	C3
0 Pence Maybe I'll wait on that.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
1 David says, I also think if you were a person who really likes doing all sorts of different things, you wouldn't be happiness.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
2 <del>Specialist role.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
3 Yeah, people are different or personalities.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
4 Some people like to have things mixed up all the time where they get bored and other people find that one thing they like and never tire of it.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
5 <del>And it really depends.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
6 <del>It really depends.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
7 I think that's why you have to take advice with a grain of salt like.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
8 Definitely, I think it's important to seek out advice from professional artists and from people that have sort of gone down the path that you're trying to go down.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
9 That's going to make things easier for you, but you know, take the advice with a grain of salt because.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
10 What might work for them isn't necessarily right right for you for any number of reasons, like what David said just from person to person, those preferences can be really different.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
11 Exactly Pablo is at the end of the day.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
12 It's about what you like doing.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
13 Alright, I'm gonna grab a little more textured brush again.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
14 I'm kind of going back and forth with harder brushes and textured brushes.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
15 I really works good to see you.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
16 <del>It's not about that.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
17 A specialist most of the time is someone with more experience.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
18 That's why he's a specialist.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
19 Oh, I understand your thoughts on that.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
20 I would agree for the most part, but I guess the part I would disagree on is that.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
21 Or what I meant is that some people start off always drawing characters, for example like they just from.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
22 I mean I like for example, I started drawing in like when I was a kid and when I was in middle school and high school was drawn.	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
23 All my favorite characters from Sailor Moon and all that kind of stuff.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
24 An anyway, some artists just like for example, love drawing characters and always draw them and never get into environments, and they go on and become a good working concepts are it could work in visual development as a character artist who knows what.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
25 Just because someone specializes, it doesn't necessarily mean that they like tried all the other things first.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
26 <del>Maybe they did.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
27 It's possible they did, but I don't think that it has to be that you have to do everything before you specialize.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
28 Some people sort of specialized the whole time just because that was their preference, you know.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
29 <del>From the beginning.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
30 If anybody is just coming in.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
31 Mohammed roller, I can show you what we did so far.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
32 This was the first study.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
33 We're doing cloud studies for part of the art club were doing clouds in different scenarios, this time different.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
34 Like lightings, different times of day.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
35 This was the second one.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
36 <del>And.</del>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
37 This is what we're working on right now.	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
38 See I do think people do though Omar, like maybe maybe you're not relating to it 'cause it's just not what your preferences are but.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
39 I for example, prefer environment designed the character design, and even though I like looking at character art by prefer to make environments, and that's not.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
40 You know, that's just my preference.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
41 All thanks claires you thank you.	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Table 17: A snippet from *Digital Painting Studies with Maddy Bellwoar-Clouds*. We show the most prominent latent codes and their representative utterances ('X'). Human annotated summary utterances are colored gray and ultra-short utterances are crossed out.