



## 2023 Design Studio

### Personas

#### Setting the scene:

The Kobian Leadership Office (KLO) at the Big Government Intelligence Agency (BGIA) is responsible for producing reliable and robust intelligence that accurately conveys the strategic plans and intentions of the leadership of a fictitious adversarial foreign nation called Kobia (where we are pretending the Richard Nixon administration is the Kobian leadership).

There are three language analysts who work in the KLO: Sloane, Cameron, and Ferris. These three analysts are responsible for processing speech-based communications in the foreign language Kobian into a form that enables their non-Kobian speaking colleagues to incorporate them into intelligence reports that draw on various sources. This processing requires the KLO language analysts to find and identify communications, or “cuts” (a single communication event or single recorded communication is often referred to as a “cut”) that are relevant to understanding Kobian leadership’s strategic plans and intentions, translate the relevant communications into English, use their subject matter expertise about Kobia and Kobian leadership to provide greater context to the communications, and help reporters piece together “stories” or reports (narratives based on data that usually comes from multiple communication events or other sources), and help prioritize which “stories” or reports are most important to get to consumers of BGIA intelligence reports.

Each of the three language analysts has their own core function in the production line:

1. Scanner/Sloane: Identifies information or activity that answers all or part of an intelligence requirement and meets reporting thresholds.
2. Translator\*/Cameron: Translate foreign language material into coherent English
3. QC-er (Quality Controller)/Ferris: Conduct quality control (QC) review on summaries and translations, and the reports that reference them.

\*Note that “translator” and “transcriber” may be used interchangeably, although in practice there is a subtle difference— “transcriber” implies the analyst may write out (transcribe) the communication in the original/native language prior to translating or reporting, whereas “translator” may indicate the analyst does not produce a native language transcription prior to producing an English language translation. Some analysts prefer or are required to do transcription prior to a translation.

The general workflow for the team is scan, translate, QC, report. The language analysts’ main responsibilities are in the scan, translate, and QC portions of the workflow. The team uses a



workflow tool to track which communications or cuts need to be processed, and which stage each communication is at:

1. Needs translation (after the scanner identifies something as relevant to an information need)
2. Needs QC (after the translator completes a translation)
3. QC'd/needs reporting (after the QC-er completes quality control)
4. Reported (after the reporter completes the report)

### More Context and Data Examples for the Nixon Administration Use Case

This [short video](#) (less than 4.5 minutes) gives a good idea of data, a data exploration tool, and type of task related to the Nixon administration use case. It gives a brief window into what a language analyst may do with one example of a tool and some types of data. It is also important as an overview of how to use the Common Analytic Platform (CAP) tool, which will be used in the Design Studio Launch Activity.

Sample data related to the Nixon use case:

- Example of an [audio recording](#) (for example from a phone conversation or a recording of a whole room) that a language analyst might hear related to the Nixon administration (imagine this is in a foreign language). As you can see from the video, there may also be associated machine-generated speech-to-text output which is imperfect (imagine that too would be in a foreign language)
- Example of a [translation](#) (a transcript in English of original foreign language material).
  - If you expand the “plus” signs in this example, you will see notes that are akin to what we call “analytic comments” or “operator comments” designed to give a reader more context than what is only in the verbatim transcript.
  - While this example shows what looks to be a verbatim translation of an audio conversation, sometimes “gists” (summaries of some portions of the data) are also done when a foreign language data file may contain sections of intelligence value but a verbatim translation of the entire file is unnecessary.
- Example of [“scan notes”](#). This is like text-based data produced by analysts **in English**, as informal notes to attach to foreign language data files when they want a quick way to document what the data is about while “scanning” or triaging data. This is in contrast to a formally produced translation which is usually done only after a file has been identified as containing intelligence value (see below).
- Examples of [documents](#) (in this case, the Presidential Daily Diary) that a language analyst might see related to the Nixon administration (imagine this text is in a foreign language). Something like this might be translated into English if it contains information of intelligence value. (Note that in our particular scenarios, the language analysts are responsible only for speech communications, not written communications/documents.)



- Example of an [intelligence report](#). LAs may also be reading intelligence reports produced by other analysts and other organizations (in addition to information on the open internet) to understand more about their topic and context. LAs do translations of foreign language data which may also be listed as a source in intelligence reporting.

## Personas

### Scanner

- **Persona:** Sloane (she/her), the scanner, is very comfortable with technology, a veteran to the KLO, and considered a “full performance” analyst. A full performance analyst has a full understanding of their field (in this case, language analysis), and can independently handle situations or assignments with minimal instruction. A full performance analyst exercises independent judgment to determine how to accomplish their work within established priorities and deadlines.
- **Job description:** The scanner’s primary responsibility is to perform content triage of very large amounts of speech communications in the Kobian language and identify information that answers intelligence requirements about Kobian leadership’s plans and intentions. Despite her veteran status, she can never be sure what data she may hear on any given day; she leverages her subject matter expertise and knowledge of current Kobian affairs to guide her exploration of the information available to her each day, but coming across new or unexpected information can challenge her assumptions about what story lines she should prioritize and how she should go about finding information, as well as her established procedures for sensemaking.
- **A day in the life:** The scanner often begins her day listening to Nixon’s conversations with his most trusted aides, as in her experience this is the best way to help her get a sense of the most up to date “big picture” on Kobian leadership plans and intentions, and she can then prioritize which story lines to pursue, based on intelligence needs that various story lines may answer. She may “gist” or summarize a few conversations that don’t contain enough significant information to warrant further processing, but are important to document in case they are needed in the future to fully understand future communications; similarly, she may update BGIA corporate knowledge bases with straightforward biographical facts revealed during her content triage. These gists or updates to knowledge bases are available for others in her office to search, but not necessarily pushed to them.

Beginning her day as she typically does, she may uncover information that leads her down a different path. Today, she finds a conversation between President Nixon and one



of his most trusted advisers about the leader of Zoolandia, one of Kobia's allied countries, visiting Kobia that day. Whatever is discussed in this meeting will be very important for understanding Nixon's plans and intentions.

There is not much more to go on, but the scanner knows that when Nixon hosts visitors in the Oval Office, he often shows them his prize music box that plays "Hail to the Chief." Because she hopes to learn more about Kobia's foreign policy intentions from listening to Nixon's Oval Office meeting with the leader of Zoolandia, she searches for communications that have the "Hail to the Chief" music in them. The first two results from her search that she listens to include other music playing, *not* "Hail to the Chief," so she takes the time to mark these as "bad" hits (a "hit" being when a machine learning model identifies something as of interest based on a particular criteria), as she knows it is important to feed truth labels to the ML Sound Event Detection Model to improve its performance. The third result from her search is indeed a recording of Nixon's meeting with the president of Zoolandia and showing him the music box; first, she marks this hit as "good," then continues listening to the Kobian-language communication. While showing off the music box, Nixon is explaining how tariffs affected import of the supplies needed for building the music boxes he loves so much, and the scanner realizes this is part of a bigger conversation on Nixon's plans for economic engagement with Zoolandia. She clicks through the conversation, listening to various snippets, to find out which portions of the recording encompass the discussion on economic engagement, so she can tee it up for the transcriber to translate. She also sends off a note to her counterpart in the Zoolandia Leadership Office (ZLO) to make sure she is aware that the President of Zoolandia is in Kobia.

As she sends a chat to the KLO translator, letting them know there is a cut about Nixon's plans for economic engagement with Zoolandia that they need to work on ASAP, she gets a message back from her counterpart in the ZLO. The ZLO scanner includes a gist she did yesterday; the information was so sparse that it didn't make sense yesterday, but in light of the information the KLO scanner has provided, the ZLO scanner now understands the Zoolandia President will also be meeting with someone on Nixon's National Security Team about Kobian military policy during his visit to Kobia, and informs the KLO scanner that their customers have expressed high interest in understanding any intent for Zoolandia and Kobia to cooperate militarily.

Because military cooperation between Zoolandia and Kobia is of higher interest to consumers of BGIA's intelligence reports than Kobian leadership plans and intentions, the KLO scanner knows she needs to pivot to trying to find more information about this meeting and what it might reveal about Kobia's military policy, but she is not sure where to start. She tries a query looking for any conversation where a member of the national security team is a speaker, but this search is too broad and brings back way too much irrelevant information to triage in a short time. Instead, she decides to try a new



capability that allows her to do semantic search by asking a natural language question and returning sentences that are semantically similar to her question; these sentences come from various communication events, and are produced by a speech to text model so they are often a bit unclear.

She quickly types “Are Kobian national security officials meeting with the Zoolandian President?” and receives back a list of possibly relevant sentences. She keys in on one that reads “I’ll be seeing that big guy from Zoomania today” as possibly being helpful to answering her question. First, she confirms through a separate search of wikipedia-like articles that Zoomania is the capital of Zoolandia, then she pulls up the audio communication that contains this sentence of interest in its associated speech to text output. Upon listening to the audio, she realizes that the speaker who indicates they will be meeting the Zoolandian President is female. Since there are very few female officials in Nixon’s national security team, the scanner redoes her previous query for cuts where a member of the national security team is a speaker (which previously returned too many irrelevant results to triage quickly), and filters her results on those that contain a female speaker. She quickly finds a few cuts where the same unknown female speaker gives details of interest on her planned meeting with the Zoolandian president, and sends these to the translator for further processing.

At this point, she has to leave for the day, and is concerned that she did not finish going through the conversations between Nixon and his most trusted aides and she may have missed communications she should have teed up for the translator. She asks the QC-er for some help with scanning more, then heads home.

#### Translator

- **Persona:** Cameron (they/them) is the translator and is uncomfortable with technology and is new to the KLO. Like the scanner, they are considered a full performance analyst. A full performance analyst has a full understanding of their field (in this case, language analysis), and can independently handle situations or assignments with minimal instruction. Full performance analysts exercise independent judgment to determine how to accomplish their work within established priorities and deadlines. However, Cameron’s understanding of Kobian Leadership is limited, as they have just recently completed entry level training and developmental activity on this target.
- **Job description:** The translator’s primary responsibility is to translate speech communications in the Kobian language into understandable English, which requires verbatim translation as well as insertion of appropriate analytic comments into the transcript to provide additional context when necessary. As they translate, they are also responsible for labeling data to improve machine learning models that benefit the whole



organization, and conducting content analysis and further content triage to build more robust “stories” to answer key intelligence questions.

- **Day in the life:** The KLO team of analysts uses a workflow management tool that helps organize their intelligence production process. Within this tool, the translator has a list of communications identified by the scanner which require a full translation (stage 1 in the “setting the scene” section). These can be organized by date, priority, or subject line; as the days go by, something that was high priority a few days ago may now be no longer relevant, or simply less of a priority, so deciding anew where to start is always the first step in the day for the translator. Since the translator is new to the KLO, they may consult with the scanner or QC-er to decide where to start with translations each day. Also, as the scanner identifies new communications for translation each day, the translator’s plans may need to shift. Similarly, as the translator digs deeper into the day’s work, they may need to share insights or information gaps back to the scanner to help direct the scanner’s efforts.

Today, they turn to translating the high priority communication about Nixon’s plans for economic engagement with Zoolandia. Since they haven’t worked economic issues much in the past, their grasp on economic vocabulary in Kobian is rusty, and they frequently need to turn away from their transcription environment, which is sort of like a word processor, to their online dictionary to look up various words. Additionally, they are working to improve their grasp of Kobian grammar, so they look back at their notes from last week’s mentoring sessions with a Quality Controller (QC-er), who corrected grammatical mistakes Cameron made in previous translations; they are striving not to make the same grammatical mistakes they have made in previous translations. They are also unfamiliar with Kobian and Zoolandian administration officials who deal with economic issues, so they frequently need to look up oblique references to people mentioned in the conversation to determine who they refer to, and then add notes about this to the transcript. This information is sometimes found in various corporate data storage tools, other translations completed in the past, and/or news reporting on Kobia and Zoolandia in both English and Kobian; other times, the translator has to wait until someone on the team with more experience is available to ask, either because the information they are looking for has not been documented before, or they simply don’t know where or how to look for the information they need.

As the translator works on this communication, they notice the automated speech to text transcript (which is presented in their audio player they have up on a separate monitor from their transcription environment) is quite errorful. Although they are not technically-minded, they vaguely understand that correcting the speech to text will lead to better output in the future, which will theoretically help everyone on the team, but it is difficult to balance their time to correct the speech to text *and* do the translation. It’s not like just a few words would need correction; they could spend all day doing corrections,



but then they wouldn't get any translation done. Ultimately, they decide to just move forward with the translations and ignore the errorful speech to text.

After listening multiple times to a part of the communication that was particularly difficult to understand given the poor audio quality, the translator picks out an oblique reference to the importance of vibranium, which their initial research indicates is one of Kobia's main exports. Cameron decides to look for more information about exactly how export of vibranium fits into Kobia's economic plans.

They are aware of a new capability that allows them to do semantic search by asking a natural language question and returning sentences that are semantically similar to the question, but they find it confusing as it's a very different way than how they are used to searching the data. They are more comfortable using traditional keyword search and going through entire communications (rather than sentences) one at a time, even though it is much more time consuming. They turn to their query tool and do a keyword search of automated speech to text output for the word "vibranium," but no communications are returned. They know, however, this doesn't mean there aren't more communications about vibranium. For example, even the cut they just translated is not returned, because "vibranium" was not correctly transcribed in the speech to text transcript. They try a few more keyword searches they think might turn up conversations about vibranium, including "export" and "mine." The former brings back no results, and the latter brings up way too many results to go through, including many mistranscriptions (where the word "mine" is in the speech to text, but isn't actually spoken) and many where it is used in a totally different context (that paper is mine). Eventually they decide just to add an analytic comment saying "no further information on how vibranium fits into Kobia's economic plans," submit the completed transcript to the workflow tool for QC, and give the scanner a heads up to be on the lookout for more information about how vibranium exports fits into Kobia's economic plans. They then turn to translating a different record.

#### QC-er

- **Persona:** Ferris (he/him), the quality controller (or QC-er), is a senior language analyst and subject matter expert on Kobian Leadership; he has many years of experience on the team, but is uncomfortable with technology. As a senior, the QC-er has extensive knowledge and experience. Given broad objectives, he can independently handle a wide range of complex and non-routine situations, act independently to establish priorities and deadlines, exercise judgment to pivot, and may lead special projects, teams, or initiatives.
- **Job description:** The QC-er's primary responsibility is to conduct quality control review on gists and transcripts completed by other members of the team, but he will also



perform scanner and translator roles as needed. He also serves as a mentor to the other language analyst members of the team, providing useful and relevant feedback or help as needed, and provides valuable insight to reporting analysts, ensuring reports accurately reflect the meaning context of language material referenced.

- **Day in the life:** The QC-er starts his day by reviewing/correcting a transcript the translator completed the day before. He is making changes in the transcription environment, and in addition to correcting the English translation, he sometimes types the exact Kobian words that were spoken in the cut into the transcript as “track changes” visible to the translator, to help the translator understand what was said in Kobian, and learn from their mistakes when they review the QC-d version. He knows that correcting the automatic speech to text output would also help the model improve, but the STT must be corrected in the audio presenter (not the transcription tool), and he doesn’t want to retype the same Kobian words he has already typed into another environment when he already has higher priority work to complete. He hopes he’ll have some extra time to make some STT corrections later. He also adds some more contextual notes to the transcript; some of the information these notes contain are just things he has in his head based on his expertise on Kobian Leadership, and some he finds in other corporate databases or news reporting. After completing his review of the transcript, he types a summary paragraph of the contents of the conversation that answer part of an intelligence requirement and meet reporting threshold at the top of the script, emails the updated version to the translator for review (with a few notes about how they could improve future translations both in terms of the language and analytic notes), then looks through the queue of QC’d scripts waiting to be reported (see stage 3 in the “setting the scene” section) to decide which scripts he should link or group this one with to help the reporters make a coherent narrative.

He checks his calendar and notes that he has a meeting on an initiative he is leading in 30 minutes. There is a big backlog of transcripts that need to be QC’ed before they can be reported, so he clicks through looking for one he should work on– he wants to balance finding something he can get done in 30 minutes with finding something that is high priority for getting reported. After identifying a cut to QC, he spends 10 minutes on it and then realizes there is a lot of hard to translate jargon, the audio quality is very bad, and there are many places that require analytic comments that the translator was unable to find amplifying information to fulfill– all of this means that although the transcript is short, it’s going to take longer to QC, and he doesn’t want to spend time getting in the flow only to be interrupted by the meeting. He now has 15 minutes left before the meeting, so decides to do some truth labeling/annotating of speech to text. He has recently noticed that the STT model is often mistranscribing some names of key players in the Nixon administration, so he decides to look for cuts with those names he can correct; finding these cuts takes 5 minutes of his 15 minutes, and he spends 10 minutes doing the corrections before going to his meeting.





After his meeting, the QC-er has a notification that a draft report based on the group or “bundle” he created earlier is ready for his review. He opens up the reporting tool and reads through it; he thinks part of it is maybe not quite an accurate reflection of the certainty or probability that an unidentified shipment is actually a shipment of music boxes, but it is based on something he QC’ed last week and he can’t remember, so he opens up his audio query tool, which he can use to pull up the audio and the completed transcript which is now associated with it in the tool, and re-listens to the relevant portion. After he listens again, he updates an analytic comment in the transcript that said “this unidentified shipment might be a reference to music box exports” to “this unidentified shipment is almost certainly a reference to music box exports, based on the conversations that occurred later and are included in this bundle.” Moving back to the reporting tool, he highlights the portion of text about the shipment and leaves a comment saying he updated the transcript to indicate this is almost certainly a reference to a shipment of music boxes.

As he finishes reviewing the report, he gets a message from the scanner saying she has to head out for the day, and she asks if the QC-er can help out by doing some more scanning. The QC-er agrees, and since he doesn’t have much time left in the day, he decides to try out a new technology, topic modeling, to help him decide where to focus in his scanning. The new topic modeling tool shows groups of communications that are similar, calculates keywords that are most relevant to each group, and displays these keywords as labels for each group. However, it was designed to work with clean text, and because of the errorful output from speech to text models, the labels that are displayed for audio data are often unhelpful. The QC-er wishes he could leverage a more useful topic modeling tool to help him and his team determine how to better prioritize the many hours of audio communications they must process to find information of intelligence value. Lacking this, he knows that the scanner usually starts her day listening to Nixon and his most trusted aides, so he decides to run a query looking for audio where they are the speakers, but he gets back a bunch of records where they are *not* the speakers— apparently the scanner has not been keeping up with truth marking speakers, so the QC-er spends some time on the tedious task of labeling good and bad hits. Since he’s listening to the audio to label the speakers anyway, he writes up a quick gist of these conversations in the transcription environment. None of it meets the threshold for reporting, so he is probably the only one who will ever read these gists, unless they become important later. At the end of his day, the QC-er sends the scanner a note explaining what he did to help with the scanning, then heads home ready for another day tomorrow.