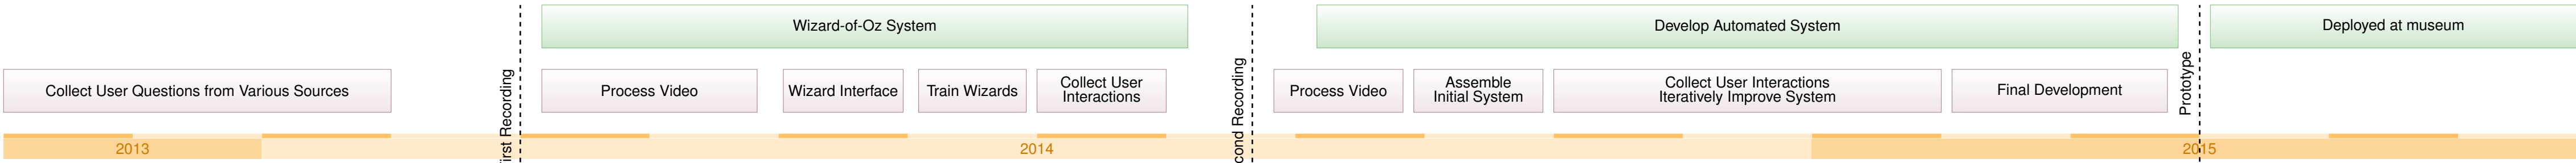
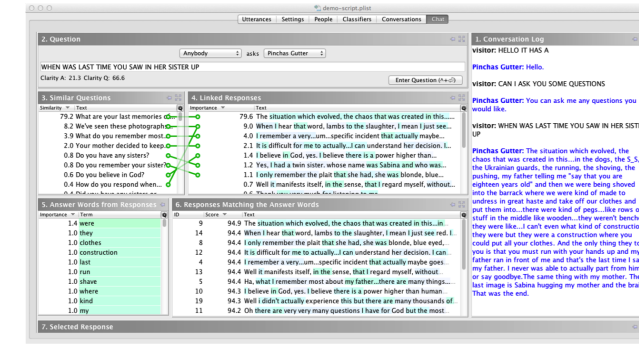



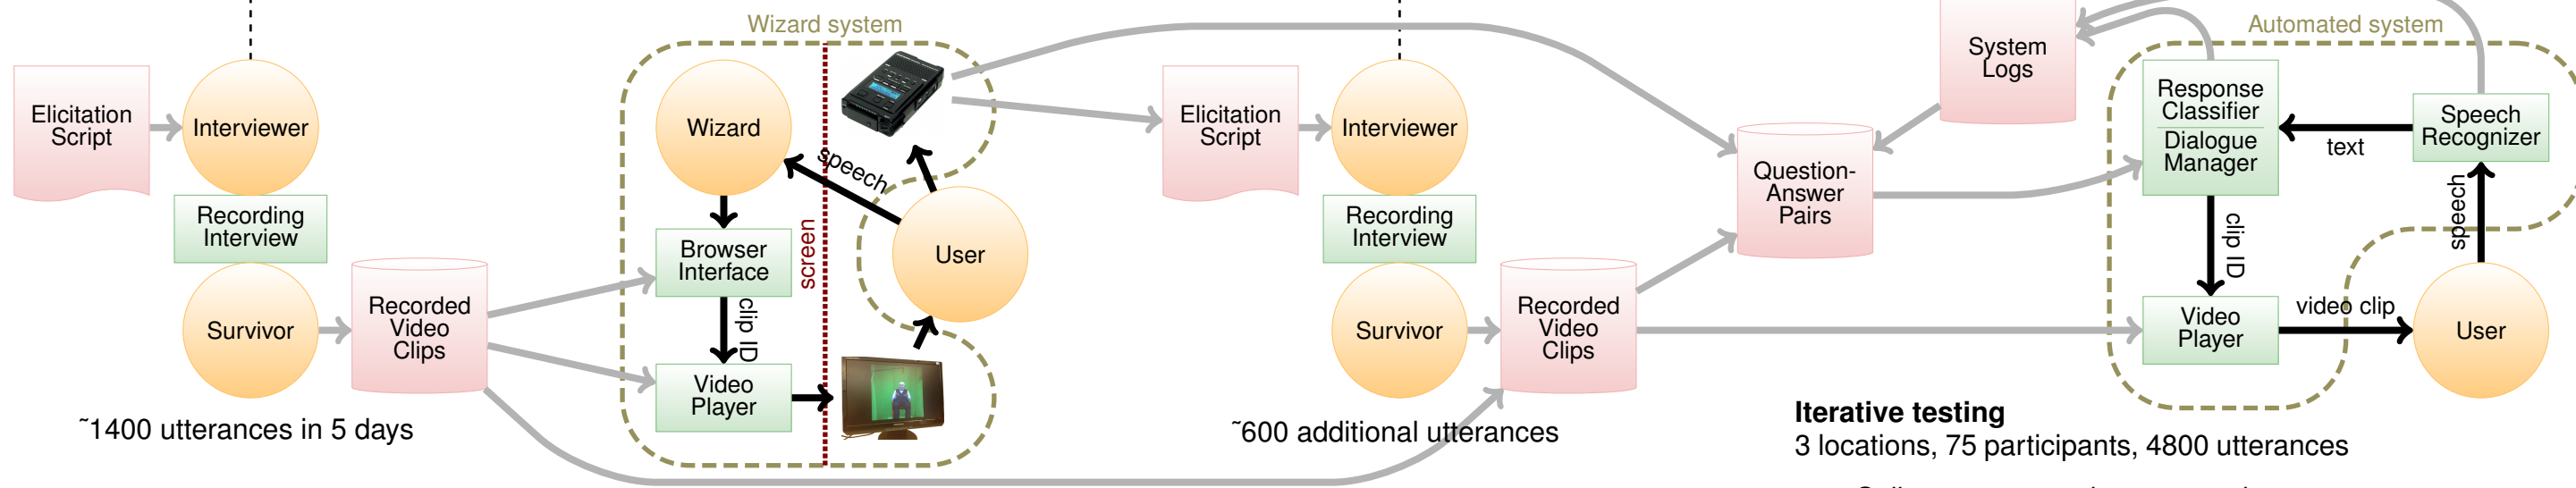

Holocaust museums: live conversations with survivors
 This generation will not be around for much longer
 New Dimensions:

- Interactive: Talk and get answers to **your** questions
- Immersive: 3D display, relighting



Common user questions

- Top 99 questions: museum experts
- New Jersey set: audience note cards
- Off-topic taxonomy



Variety of interviewers

- Stephen Smith (long relation)
- Unfamiliar adults
- Children



How many user questions are addressed by the recorded statements?

Question source	Top-99		NJ		Wizard	
	N	%	N	%	N	%
Answer available	73	74	366	49	756	58
No answer	18	18	348	47	376	29
Both	1	1	5	1	21	2
Unannotated ^a	7	7	27	4	141	11
Total	99	100	746	100	1294	100

^aUtterances left unannotated do to a technical error

How important is the second recording?

95% of user utterances have a direct response

Best Response	Dec. 10		Jan. 15		Jan. 18	
	N	%	N	%	N	%
1st rec.	452	69	427	65	203	58
2nd rec.	187	29	214	32	129	37
Off-topic	15	2	21	3	19	5
Total	654	100	662	100	351	100

Source of best response, data from system testing 2014–2015

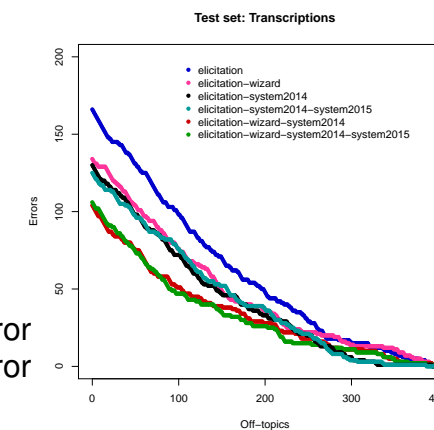
Iterative testing

3 locations, 75 participants, 4800 utterances

- Collect utterances in conversation
- Transcribe, annotate
 - Rate system response
 - Identify best response
- Retrain system

Error-return plots

20% off-topics → 16% error
 30% off-topics → 10% error



NPCEditor: classification-based language understanding

- Direct mapping from user utterances to responses.
- Selects best response from available set.
- Trained on many-to-many mapping:
 - ~1700 Pinchas utterances
 - ~10,000 user utterances
 - ~13,000 links

- Questions and responses need not be in the same language.

Cross-language information retrieval

- Estimate of the most likely language model for the response to Q:

$$P(w|Q) \cong \frac{\sum_j \phi_{R_j}(w) \prod_{q \in \text{tok}(Q)} \pi_{S_j}(q)}{\sum_j \prod_{q \in \text{tok}(Q)} \pi_{S_j}(q)}$$

Where j iterates over all linked question-response pairs.

- Rank available responses based on distance to estimated LM.
- Top-ranked response is the most appropriate.

Dialogue manager: handling misunderstanding

- Direct response not always available: *What's your favorite restaurant?*
- Is the top-ranked response close enough to the desired answer?
 - Yes: **Return top-ranked response**
 - No: **Return off-topic response**
- 5-stage off-topic policy: 1. Please repeat; 2. I don't know; 3. I can't answer; 4. Why don't you ask me about... 5. Let me tell you about...

USC Institute for Creative Technologies

Partners: USC Shoah Foundation
 Conscience Display