So...

This turned out to be pretty difficult

Where we Started

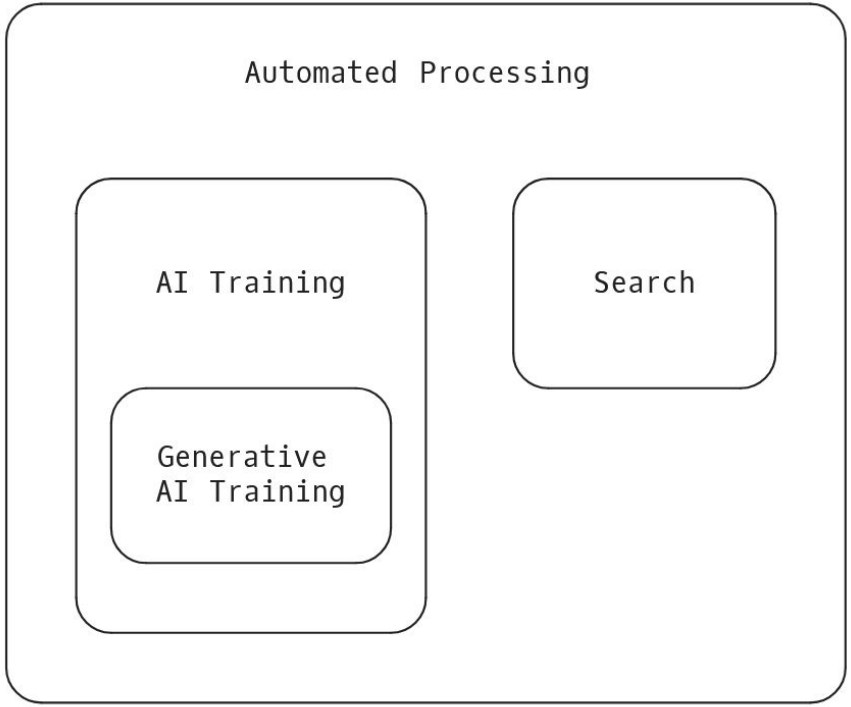

What we learned

Search isn't an agreed concept

Search uses AI (not new)

… and it uses a lot of it, including training, in operation

Search, Prior

Using one or more assets in a search application that directs users to the location from which the assets were retrieved.

Search applications can be complex and may serve multiple purposes. Only those parts of applications that direct users to the location of an asset are included in this category of use. This includes the use of titles or excerpts from assets that are used to help users select between multiple candidate options.

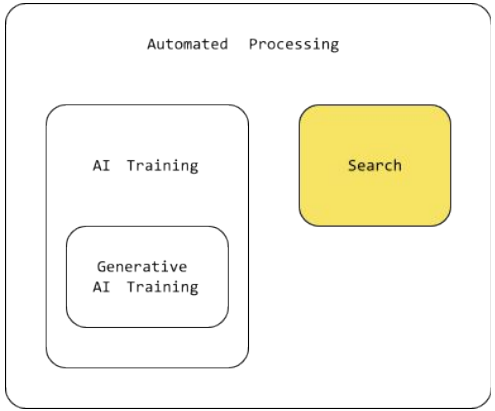

Search, Prior (...cont.)

Preferences for the Search category apply to those parts of applications that provide search capabilities, regardless of what other preferences are stated.

Parts of applications that do not direct users to the location of assets, such as summaries, are not covered by this category of use.

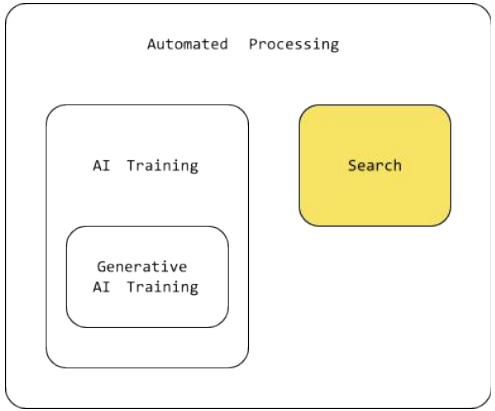

Where We Ended Up

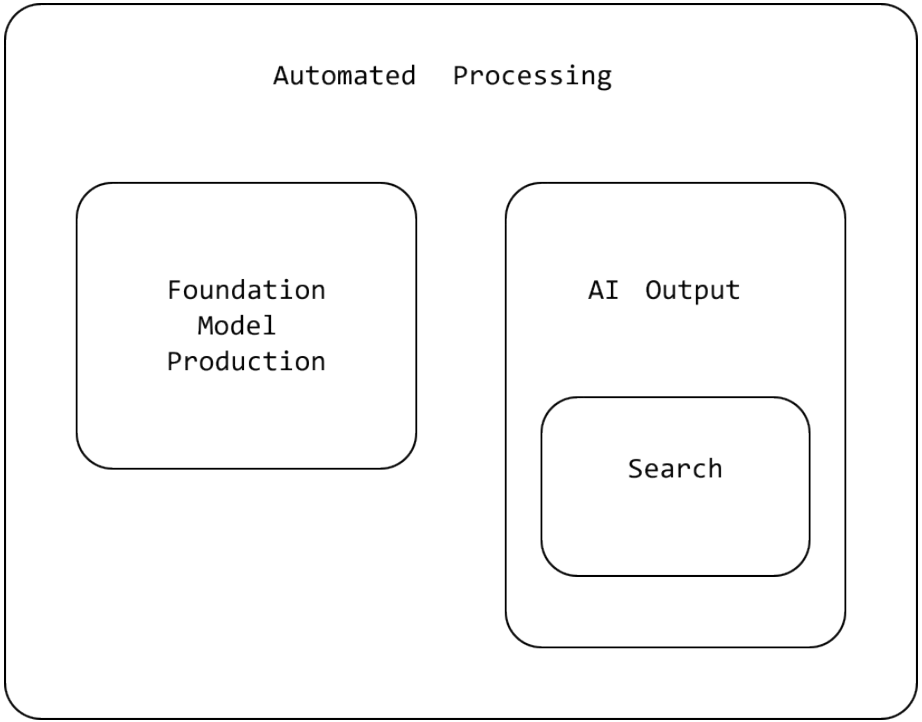

A Top-Level Category

Lots of debate at the recent interim

No resolution

Central point of tension about the unknown:

- For: cover unexpected use cases
- Against: unintended consequences

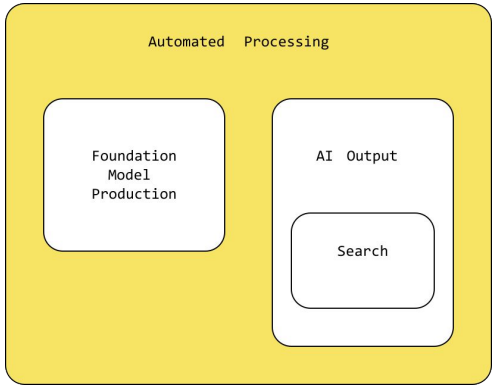

Foundation Model Creation

Making a foundation model using the asset

Successor of “generative AI training”

Premise is

- “foundation models” is a better defined term
- focus on output: a foundation model
- avoid contention in definition of “AI model”

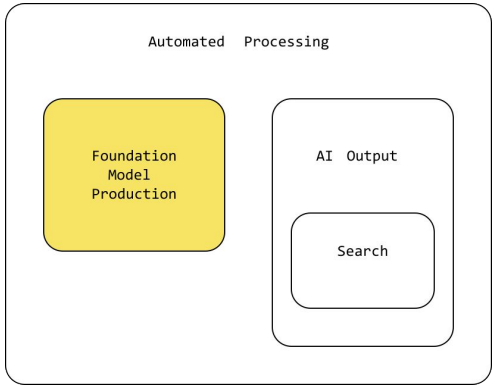

AI Output

A new concept developed during the meeting

Idea:

- Concentrate on observable outcomes
- Address substitutive concerns directly

Allow any internal processing necessary

- *including training*

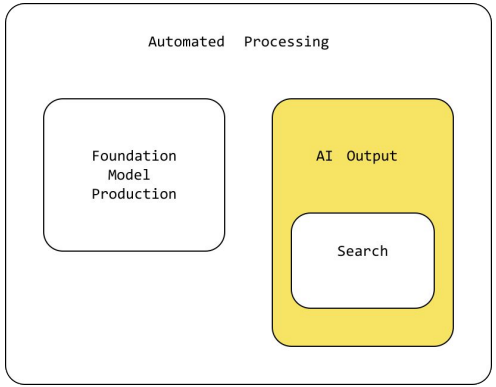

AI Output + Excerpts and Links

"Exact Text Match"

Includes internal use, *including training*

If included in outputs, need to

only include verbatim quotes, and

link back to the original source

Lots of Open Issues

We already have a lot of open issues

Becoming apparent that the “new” direction also has problems:

- Should we have a top-level category?
- What is the difference between model training categories?
- Many issues around AI Output and Search

Fine-Tuning of Foundation Models

Does the definition of foundation model training include

- fine-tuning
- Low-rate adaptation (LoRA)
- other techniques that might use assets

What does "Output" Mean?

Some suggest presentation to a human user

- Is that realizable?
- Would "at the system boundary" be more practical?

Does a preference of "n" for AI Output still allow for internal use?

- How does the system avoid having output affected by the asset?

What is the test that could be applied to determine compliance?

Excerpts and Links

What do we call this? (Pro tip: not lederhosen)

Is this rightly a subcategory of AI output?

Is this a good approximation for “search” in its most naive form?

Or is search its own category?

In England we have a saying for very situation such as this, which is that it it's ***difficult difficult lemon difficult****.*

You bet you huh. I do think a consensus is emerging.

No, no it's not and you have no basis for saying that. That is ridiculous.

Calm down because we don't want you to have another hemorrhage.