

ORACLE®

Oracle Digital Assistant

The Complete Training

Training your model

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Topic agenda

- 1 ➤ Recap – the models in Oracle Digital Assistant
- 2 ➤ General guidelines for training your skills
- 3 ➤ Specific tips for training your skills
- 4 ➤ Tips for designing intents
- 5 ➤ Training the digital assistant
- 6 ➤ Quality reports


Topic agenda

- 1 ➤ Recap – the models in Oracle Digital Assistant
- 2 ➤ General guidelines for training your skills
- 3 ➤ Specific tips for training your skills
- 4 ➤ Tips for designing intents
- 5 ➤ Training the digital assistant
- 6 ➤ Quality reports

Why have different training models?

- Trainer Ht
 - Rules based
 - Fast & best suited for small set of utterances
 - Good for new development
- Trainer Tm
 - Machine learning
 - Thrives on more and more data
 - Higher accuracy (especially data outside your utterances) if enough data
 - Already trained on “knowledge” of English language (NLP)
- Q&A

Your digital assistant is only as **smart** as the data you train it with
(rubbish in, rubbish out)

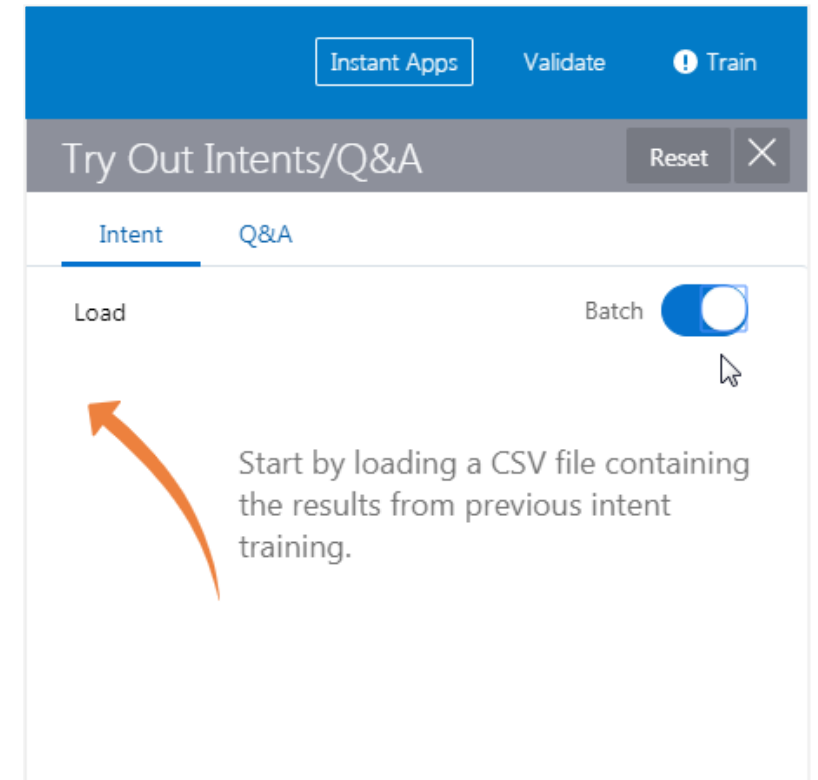


Topic agenda


- 1 ➤ Recap – the models in Oracle Digital Assistant
- 2 ➤ General guidelines for training your skills
- 3 ➤ Specific tips for training your skills
- 4 ➤ Tips for designing intents
- 5 ➤ Training the digital assistant
- 6 ➤ Quality reports

General guidelines for training your skills

- You MUST have buy in for the long term
- Your first release will have the language understanding of a 2 year old (if you are lucky!)
 - But that's ok, you need to start somewhere
- Plan for NLP improvement
 - The only real data is user data
 - Train, test, repeat
 - Use batch training feature
 - For all new data perform a 80/20 split
 - 80% training data, 20% testing data



Many bot projects **fail** because of **unrealistic expectations**. Get **buy in** for the **long term**, start **small**, grow.



General guidelines for training your skills

- Your goal is to train the model on real-world sample utterances
 - Synthesized utterances are no substitute for real world utterances
- Real-world data will likely need to be manually classified
 - It is not always clear what phrases map to an intent
 - Needs to be decide by project team: business, conversational designer, developer
- You shouldn't necessarily sanitize user data
 - Include common mis-spellings, slang, synonyms, abbreviations where it makes sense


General guidelines for training your skills

- The golden rules to help you develop good understanding
 - For production use Trainer Tm
 - Real world utterances
 - Well design and classified intents
 - Train unresolved (anti-usecase) utterances
 - As many as you need to get the results you desire
 - Our research has shown a “plateauing off” around the 92% accuracy level


Synthesizing utterances when you have no existing data

- Trainer Ht best place to start off
 - Generally better with smaller data sets
 - Plan to move to Trainer Tm as you gather more sample user inputs
- You have to synthesize sample utterance
 - Your primary goal is to help disambiguate intents
- A model has no inherent knowledge of what an utterance actually means
 - Frequency of words, sentence patterns, some knowledge of parts of speech, train with synonyms
- Train, test, repeat

So if I am going to build a **good model** shouldn't I know how it works?




Pizza




Banking

Pizza




Banking

Pizza

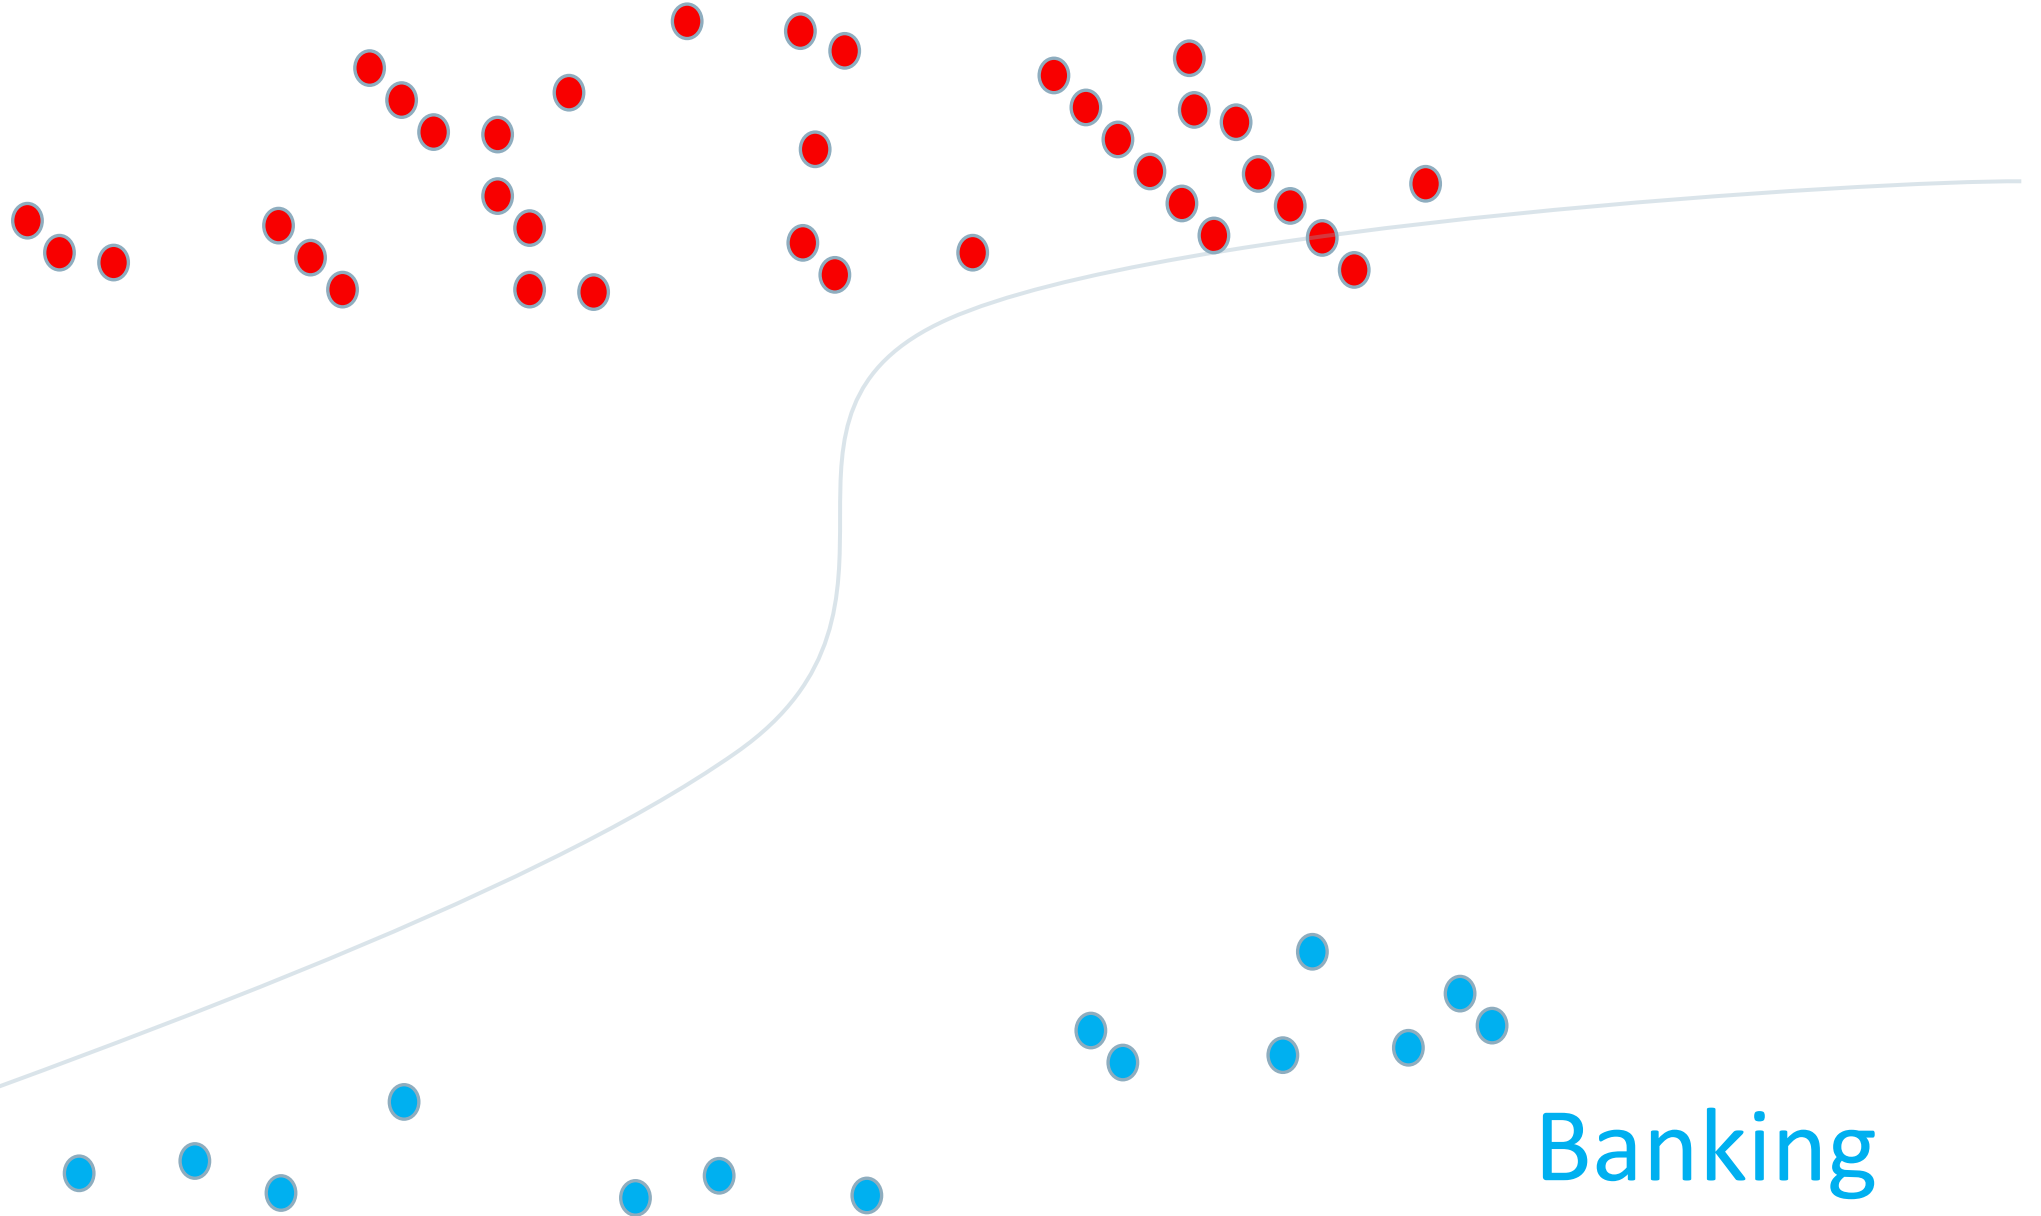


Banking

Pizza




Pizza




Banking


Pizza



Banking




Pizza




Banking

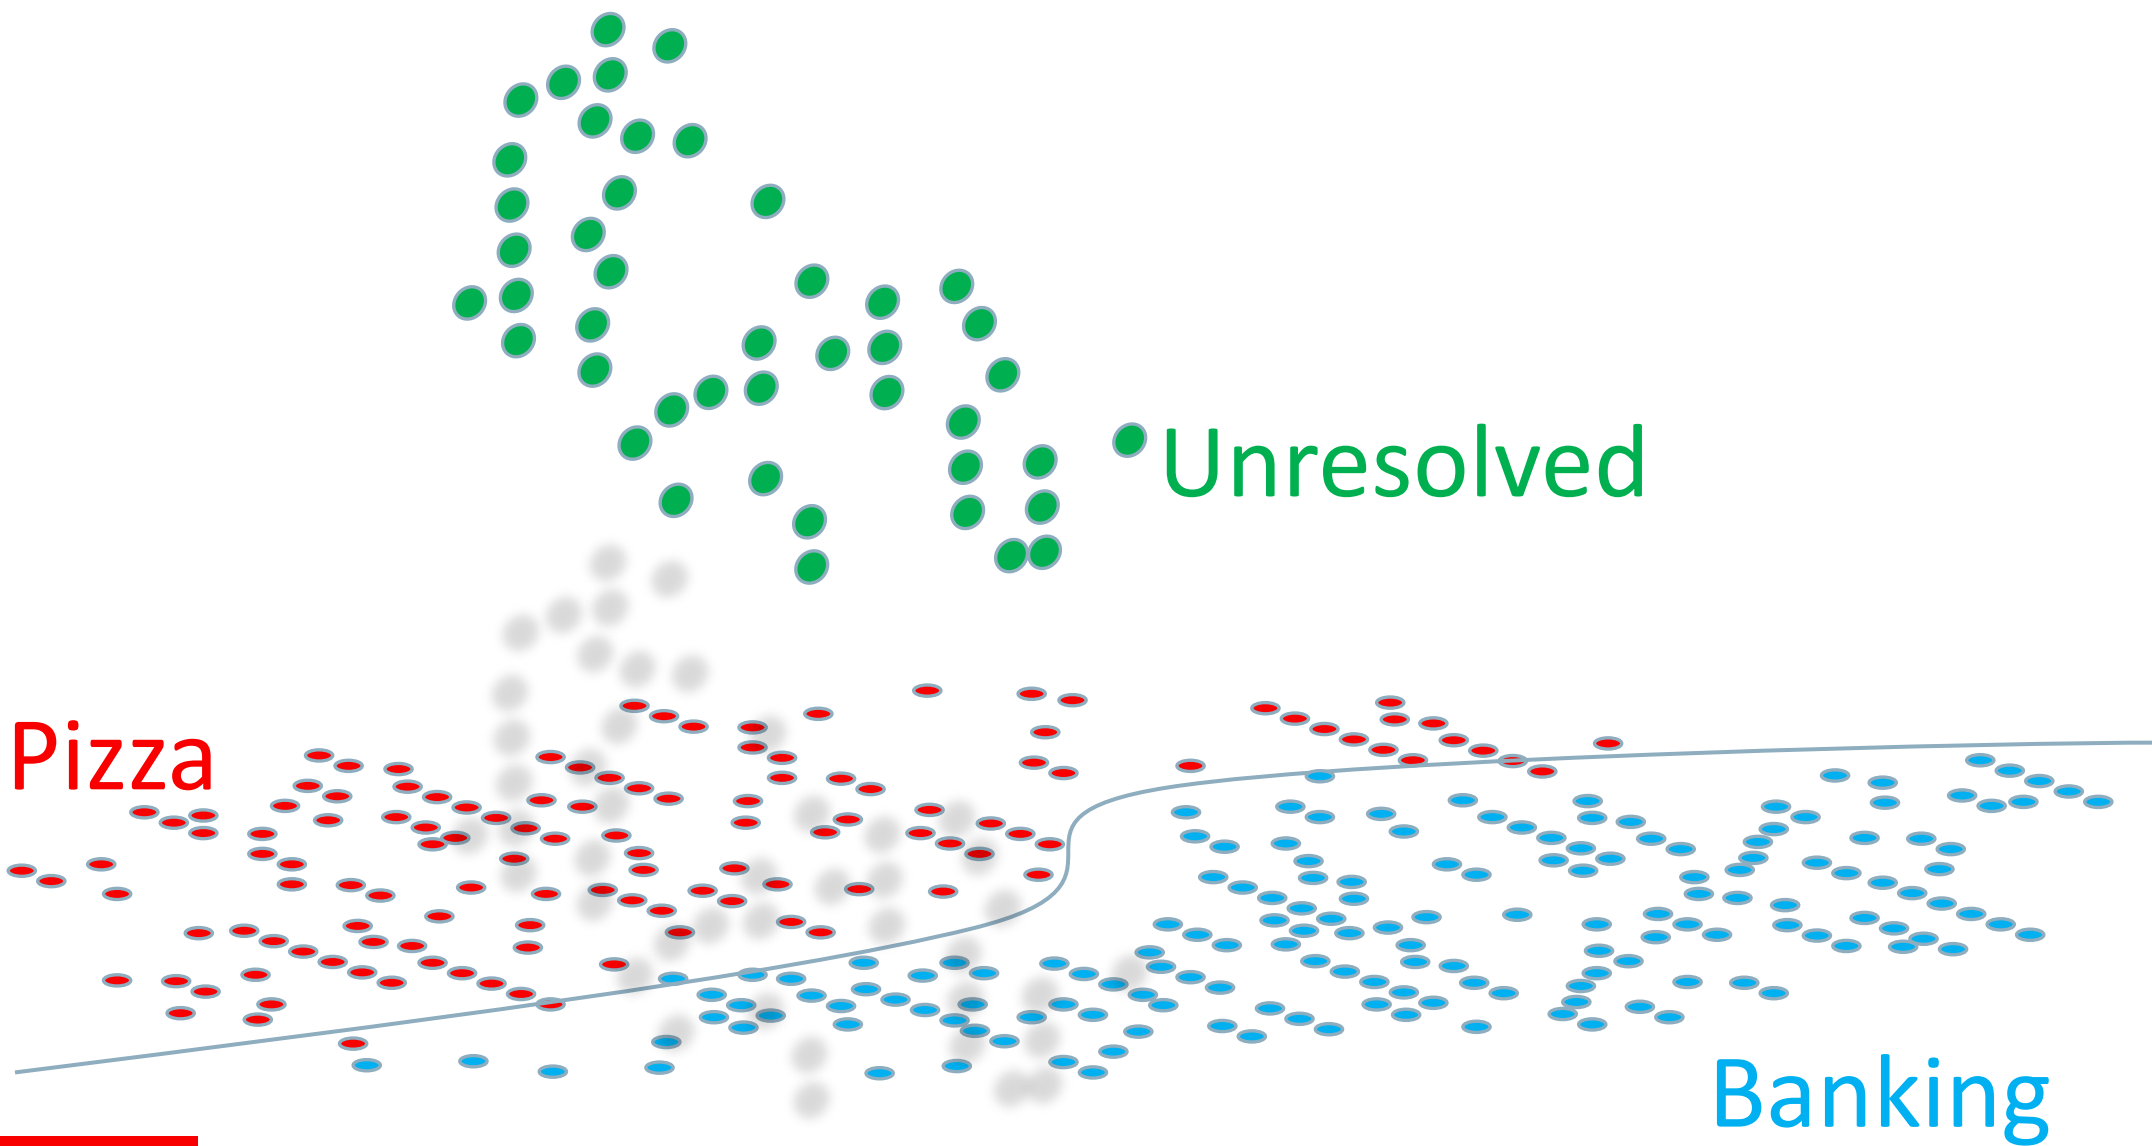
If your bot is only trained to know about **two possible intents**, then it will try to resolve any input to **one of those intents**




Pizza



Banking



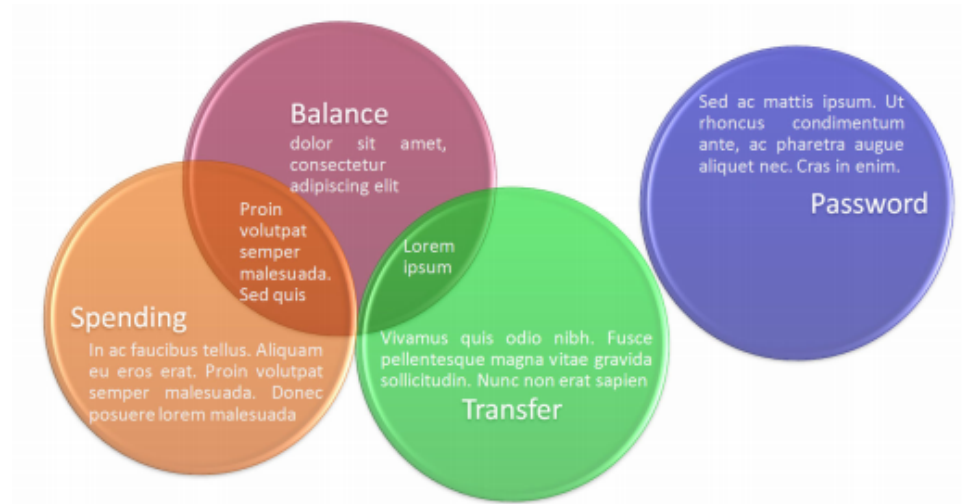


Topic agenda

- 1 ➤ Recap – the models in Oracle Digital Assistant
- 2 ➤ General guidelines for training your skills
- 3 ➤ Specific tips for training your skills**
- 4 ➤ Tips for designing intents
- 5 ➤ Training the digital assistant
- 6 ➤ Quality reports


When synthesizing utterances what makes an intent unique

- The goal is to disambiguate intents
- What are key actors, actions and objects
 - “who” does “what” to “who”
 - Design utterances around these
- Consider combining intents if too similar
 - E.g. home, travel, car insurance
- Extra words can be used to weight utterances
 - But always retest when you do this



When synthesizing utterances be careful of lexical chrome

- Some words appear across all intents but add no meaning
 - “please”, “thank you”
- Ensure you don’t unintentionally skew results
 - E.g. if one intent consistently uses “please” that can skew results if the user says “please”
- Entity value in training utterances might inadvertently skew intent resolution
 - BOOK_ORDER_INTENT “Please order All Aboard the Toilet Train”
 - Incorrectly resolves “Does the train have a toilet”



When synthesizing utterances consider spelling & grammar

- Don't have to cover every form a word (Trainer Ht)
 - Different forms of a word are reduced to a common root
 - “want”, “wanted” / “run” “ran” “running” “runs”
- Sentence structure may aid intent resolution (Trainer Ht)
 - Therefore use grammatically correct sentence where possible
- However, your data should still reflect your users' language
 - “wanna”, “gonna”, “gotta”, “dunno”

When you see commonality consider combining intents

- More commonality of terms between intent utterances, the more difficult to disambiguate intents
- Consider creating one common intent and use entities to understand the discriminating factors

Updating Policy to add new person

Add wife to my policy

Alter my insurance to add son

Can I change my policy to include spouse

Add my wife to insurance

Alter my policy to add wife

Update policy to include son

Updating Policy to protect no claims

Add no claims to my policy

Alter my insurance to add protection

Can I change my policy to protect NCB

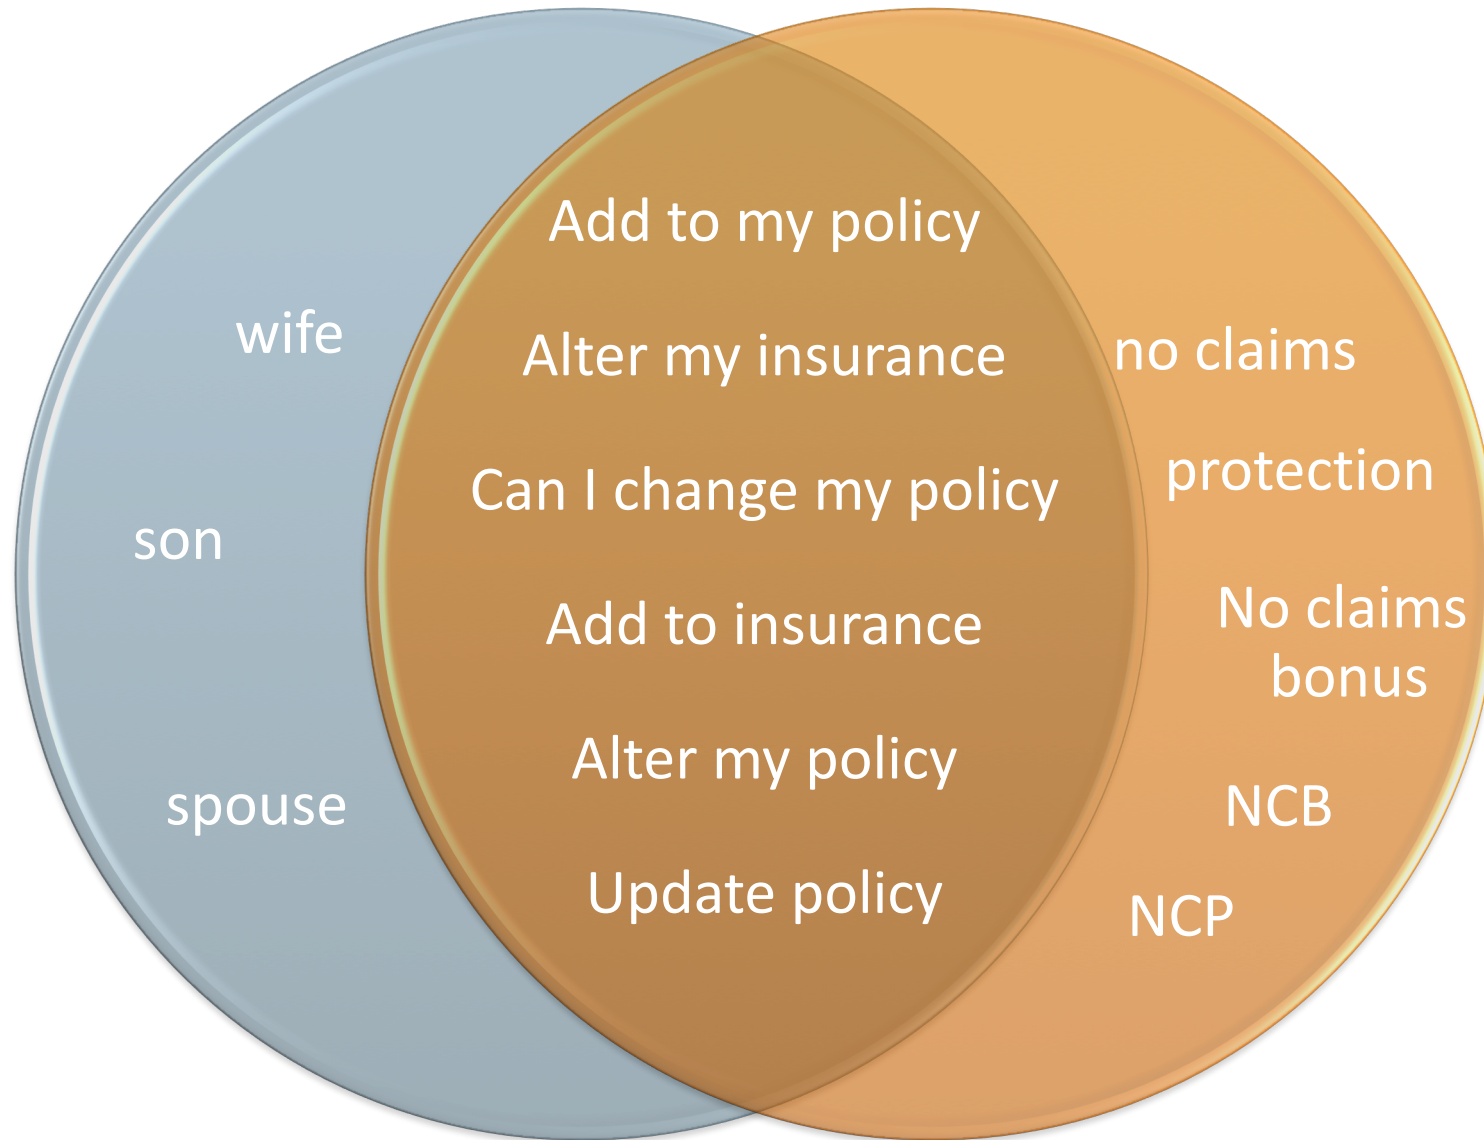
Add no claims protection to insurance

Alter my policy to add no claims bonus

Update policy to include NCP

Updating Policy to add new person

Updating Policy to protect no claims



Updating Policy

Person Entity

- spouse
- wife
- son

Add to my policy

Alter my insurance

Can I change my policy

Add to insurance

Alter my policy

Update policy


Policy Claim Entity

- No claims/
protection
- No claims bonus
- NCB/NCP

Topic agenda

- 1 ➤ Recap – the models in Oracle Digital Assistant
- 2 ➤ General guidelines for training your skills
- 3 ➤ Specific tips for training your skills
- 4 ➤ Tips for designing intents**
- 5 ➤ Training the digital assistant
- 6 ➤ Quality reports

Databases are **designed**, cars are designed, UX is **designed**. And so **Intents** must also be **designed**



General tips for designing intents

- There are no hard and fast rules but:
 - 1 intent is too few, 500 intents per skill is probably too many
 - Don't overuse skills simply to partition, but don't underuse them either
- Each intent should have a distinct use case
 - If you see commonality consider refactoring
- Be “smart” when you train your intents
 - Imbalanced training can cause imbalances results
 - Rubbish in rubbish out

Good practice for designing intents

- Create intents for things you know skill can't handle but likely to be asked
 - Smalltalk, swearing, common business uses cases not handled by the skill
- A single use case might be implemented by multiple intents
 - Returns policy might be handled by 3 different intents (all execute the same flow)
 - Potentially cleaner classification
 - Opportunity to handle each intent differently in the future
- Consider intent negation – but this is a problem not easily solved
 - "I want to continue my subscription".
 - "I don't want to continue my subscription"
 - "I want to discontinue my subscription"

Train the unresolved intent

- Capture phrases which your bot should consider unresolved
 - Create unresolvedIntent intent
 - Use this to record any phrases you DON'T want to be resolved to your actual intents
- Analyze bot conversations and train the bot with any malicious/mismatched/rogue input
- Gives the intent resolution engine knowledge of what is NOT an intent

Topic agenda

- 1 ➤ Recap – the models in Oracle Digital Assistant
- 2 ➤ General guidelines for training your skills
- 3 ➤ Specific tips for training your skills
- 4 ➤ Tips for designing intents
- 5 ➤ Training the digital assistant**
- 6 ➤ Quality reports

Training the digital assistant

- All training guidelines apply to the skill
- DA creates “uber model” – an intent for each skill
 - This is only helps DA to identify candidate skills
 - The more skills the lower you should consider threshold
- There are only three intents at the digital assistant level
 - Help, Exit, unresolvedIntent for disambiguating help and exit only.
- Digital assistant treats a skill unresolvedIntent in a special way
 - If no active skill, shows all available skills in the DA automatically
 - If a skill is active, the skill’s unresolvedIntent handles the user input
 - You should never be prompted with the unresolvedIntent

Candidate Skills Confidence Threshold

0.4





Topic agenda




- 1 ➤ Recap – the models in Oracle Digital Assistant
- 2 ➤ General guidelines for training your skills
- 3 ➤ Specific tips for training your skills
- 4 ➤ Tips for designing intents
- 5 ➤ Training the digital assistant
- 6 ➤ Quality reports

Introduction to quality reports

- Report that highlights possible intent resolution errors
 - Where utterances within intents are too similar
- Use a random 20% of the utterances to test the other 80%
- Compares all possible combination of pairs of intents in order to report
 - High quality: intents are distinct
 - Medium quality: some similar utterances
 - Low quality: the intent pairs aren't differentiated enough
- Indicates the number of utterances which may be problematic

 Int_Agnt_Exm_Authrty
↔ Int_Agnt_Min_Age

 Int_Agnt_Shift_Agncy
↔ Int_Agnt_Shift_Cmpny

  Int_Agnt_Min_Age
↔ Int_Agnt_Rtrmnt_Age 

Quality report on utterances

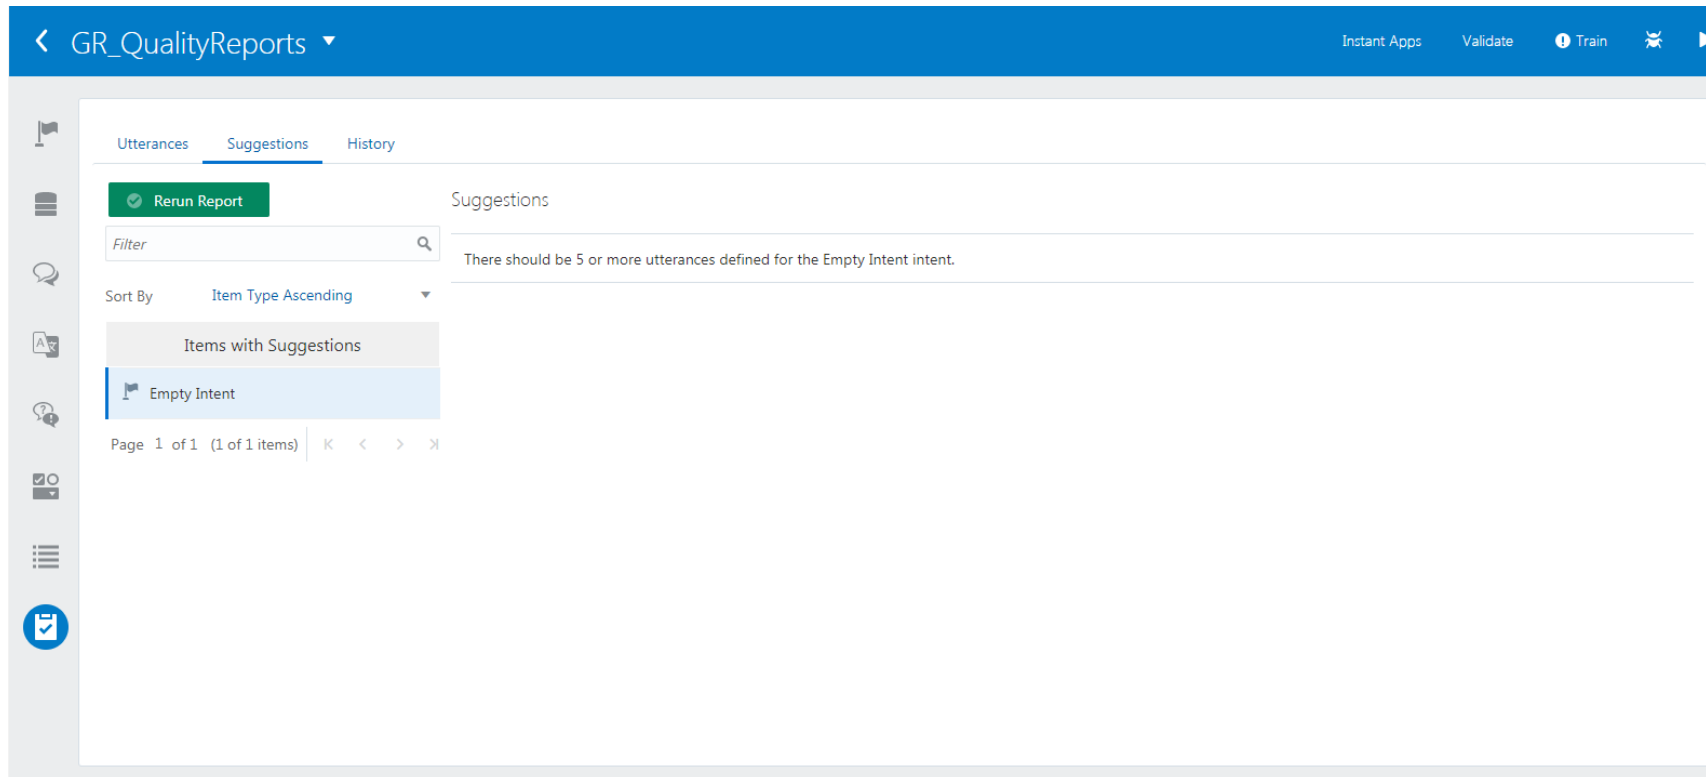
- Indicates an utterances that a deemed to be similar
- Misclassified utterances (tests did not resolve to correct intent)

The screenshot displays the Oracle Quality Reports interface. The top navigation bar includes 'GR_QualityReports', 'Instant Apps', 'Validate', 'Train', and a search icon. The main content area is divided into 'Utterances', 'Suggestions', and 'History' tabs. A 'Rerun Report' button and a 'Show All' toggle are visible. A search filter is present, and the results are sorted by 'Quality Ascending'. The interface shows two main sections: 'Similar Utterances' and 'Misclassified Utterances'. The 'Similar Utterances' section shows a pair of intents: 'Int_Agnt_Monthly_Statmnt' and 'Int_Agnt_Periodicity_Paymnt', with example utterances 'When will i get monthly commission statement' and 'When will i get commission'. The 'Misclassified Utterances' section shows a table with columns for 'Utterance', 'Expected Intent', 'Observed Intent', and 'Accuracy'. The table lists three misclassified utterances with their respective expected and observed intents and accuracy scores (indicated by red exclamation marks).

Utterance	Expected Intent	Observed Intent	Accuracy
When will i get commission	Int_Agnt_Periodicity_Paymnt	Int_Agnt_Monthly_Statmnt	!
Will i get my commission every month	Int_Agnt_Periodicity_Paymnt	Int_Agnt_Monthly_Statmnt	!
When will i get monthly commission statement	Int_Agnt_Monthly_Statmnt	Int_Agnt_Periodicity_Paymnt	!

Quality report suggestions

- Reports any suggested changes you can make to your intents or utterances



The screenshot displays the Oracle Quality Reports interface. At the top, a blue header bar contains the breadcrumb 'GR_QualityReports' and navigation links for 'Instant Apps', 'Validate', 'Train', and a play button. Below the header, the interface is divided into three tabs: 'Utterances', 'Suggestions', and 'History'. The 'Suggestions' tab is active. A green 'Rerun Report' button is visible. A search filter is present, and the results are sorted by 'Item Type Ascending'. A table titled 'Items with Suggestions' shows one item, 'Empty Intent', which is highlighted. Below the table, pagination information indicates 'Page 1 of 1 (1 of 1 items)'. To the right of the table, a message states: 'There should be 5 or more utterances defined for the Empty Intent intent.'

Quality report history

- View past user input ranked by win margin and confidence level
- Useful for
 - Narrow margins between intent classification
 - Intent resolution failures
 - Low confidence resolution

The screenshot displays the Oracle Quality Reports interface. The top navigation bar includes 'GR_QualityReports', 'Instant Apps', 'Validate', 'Train', and a play button. The main content area is divided into 'Utterances', 'Suggestions', and 'History' tabs. A search filter is applied: 'Show me all customer messages' from 'Last 30 Days' where 'All' of the following are true: 'Win Margin' is 'Is Less Than' '10%'. Below the filter, a list of customer messages is shown, each with its corresponding intent data and confidence level. The first message is 'what do i do once I passed test' with intent 'Int_Agnt_Afr_Exam_Procedure' and a confidence of 18%. The second message is 'where can I sit an exam' with intent 'Int_Agnt_Passing_Marks' and a confidence of 14%. The third message is 'Is there a method for applying to be a specialist on the web'. The fourth message is 'can I add my wife'. The fifth message is 'how much to be an agent'. The interface also includes an 'Export' button, a 'Search' button, and an 'Add Example' button.

Customer Message	Intent	Win Margin	Confidence
what do i do once I passed test	Int_Agnt_Afr_Exam_Procedure	3.9%	18%
where can I sit an exam	Int_Agnt_Passing_Marks	N/A	14%
Is there a method for applying to be a specialist on the web			
can I add my wife			
how much to be an agent			

Quality report history

- Show top intents that resolved with low confidence

Top Intent Confidence ▼ Is Less Than ▼ 50% ▼ ^ ×

- Show where there may have been ambiguity between intents

Win Margin ▼ Is Less Than ▼ 20% ▼ ^ ×

- Show me used input where the top intent was unresolved

Top Intent Name ▼ Is Equal To ▼ unresolvedIntent ×



Oracle Digital Assistant Hands-On

TBD

ORACLE®