

Conversational commerce:

Emerging architectures for smart
& useful voice and chatbots

Victoria Livschitz
Founder and CTO

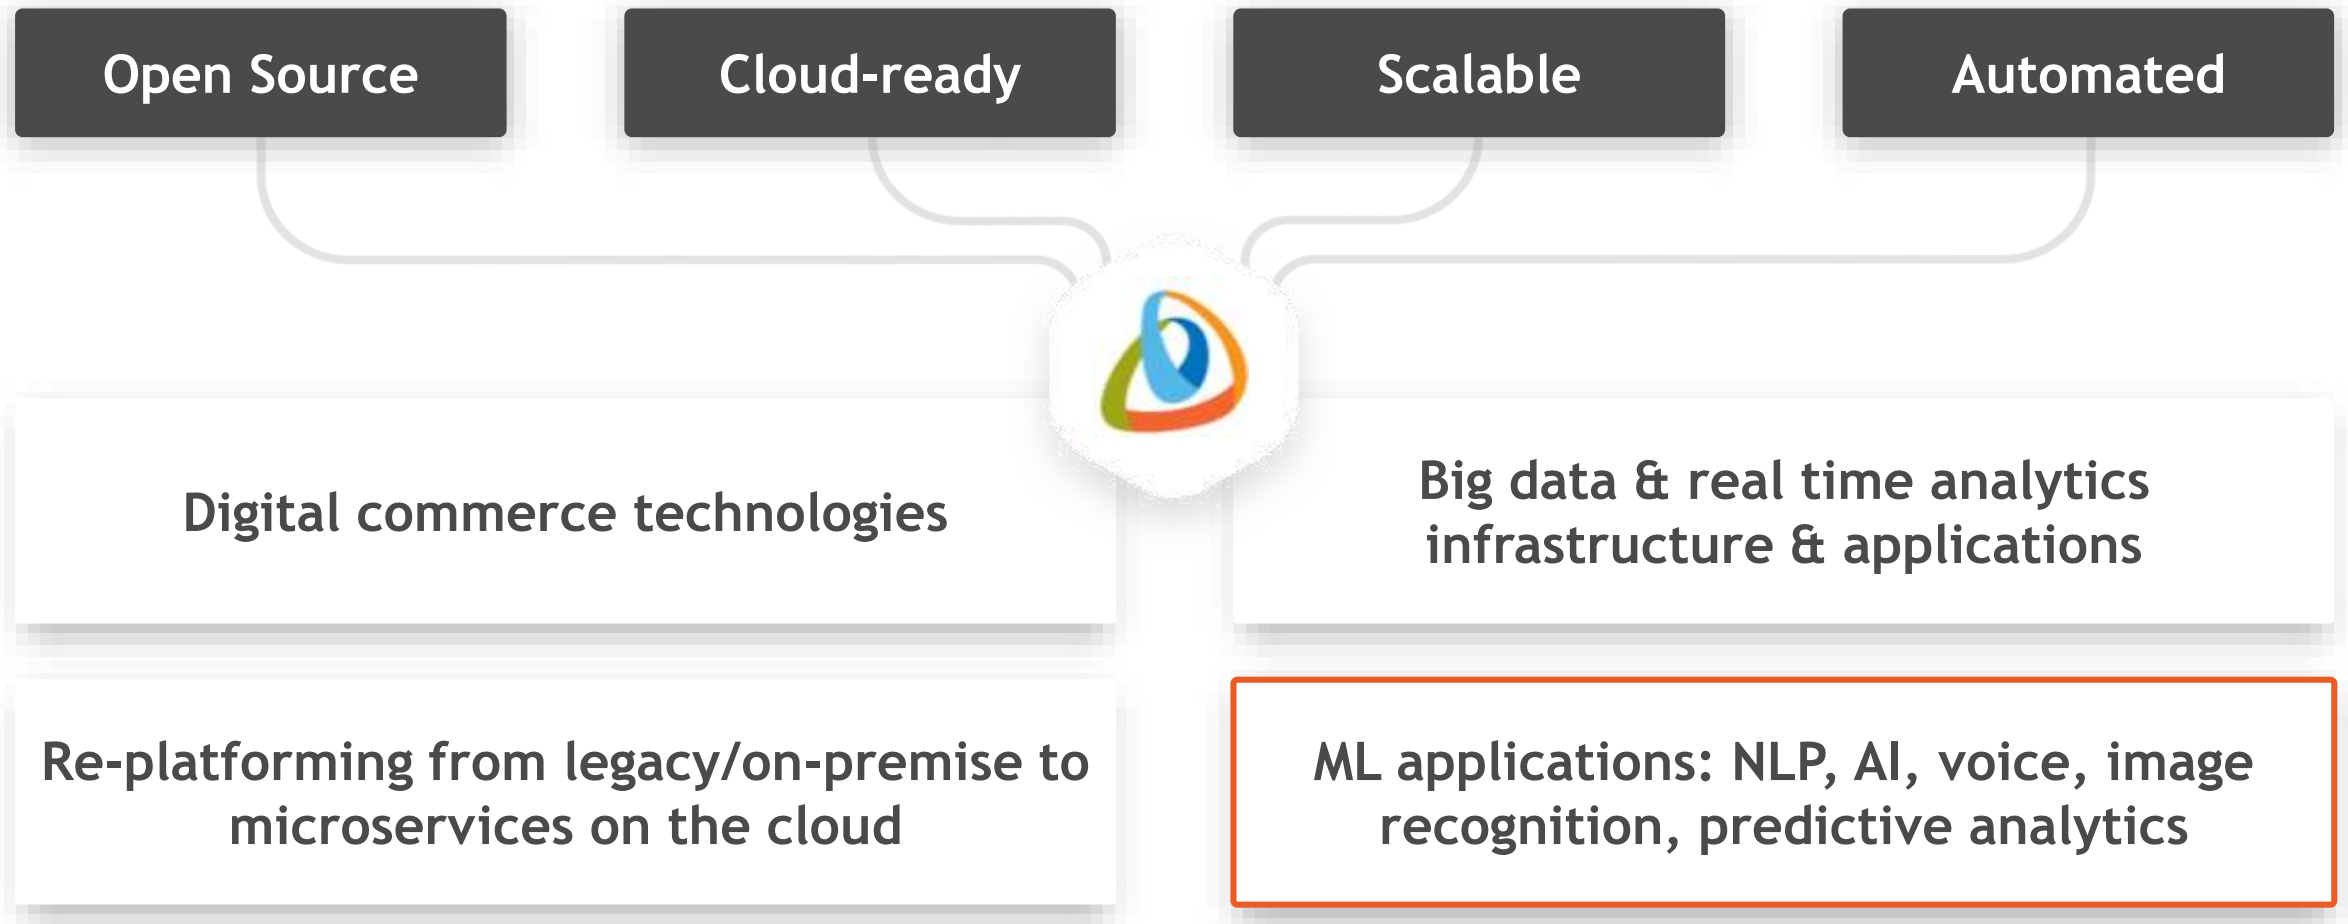


About the speaker

- Founder & CEO of Grid Dynamics from 2006, transition to CTO role in 2015
- Principal engineer at Sun; lead architect of SunGrid, world's first public cloud, 1997-2006
- HPC engineer at Ford Labs 1994 - 1997



About the company: experts in digital transformation thru emerging technologies



Thesis: voice commerce is the next big thing after mobile

Voice devices entering households

- 9M smart voice devices shipped in Q1 2018
- 50% household penetration expected by 2022
- IOT makes voice control an expected feature

Technology advances by leaps and bounds

- Speech2text approaches human quality
- Deep learning-based NLU / NER enters mainstream
- Plethora of AI platforms on the market

Massive investment

- From tech giants - over \$1B IBM Watson, etc.
- From VC - over \$700M in 2017
- 20 AI acquisitions - siri, api.ai, viv

Better interface for many digital interactions

- Reduces engagement barrier
- Perfect for search, Q&A
- Interactive clarification of intent

Types of CUI experience

Conversational commerce:

- Deep, branched conversation
- Discovery, selection, recommendation
- Large result sets, ambiguity, comparisons
- Deep integration with search, catalog, ontology
- Customer-trained NLU
- Smart dialog management to arbitrate and orchestrate multiple models

Other “deep” conversational applications

Infobot:

- Fact-oriented
- Few intent-reply patterns

Virtual assistant:

- Task-oriented
- Few intent-entity-reply patterns

Shallow, linear conversation

Smalltalk: entertainment, empathy, human mimicry

Types of CUI experience

Conversational commerce:

- Deep, branched conversation
- Discovery, selection, recommendation
- Large result sets, ambiguity, comparisons
- Deep integration with search, catalog, ontology
- Customer-trained NLU
- Smart dialog management to arbitrate and orchestrate multiple models

Other “deep” conversational applications

Infobot:

- Fact-oriented
- Few intent-reply patterns

Virtual assistant:


- Task-oriented
- Few intent-entity-reply patterns

Shallow, linear conversation

Smalltalk: entertainment, empathy, human mimicry

Grid Genie: Conversational demo by Grid Dynamics

R&D project from Grid Labs



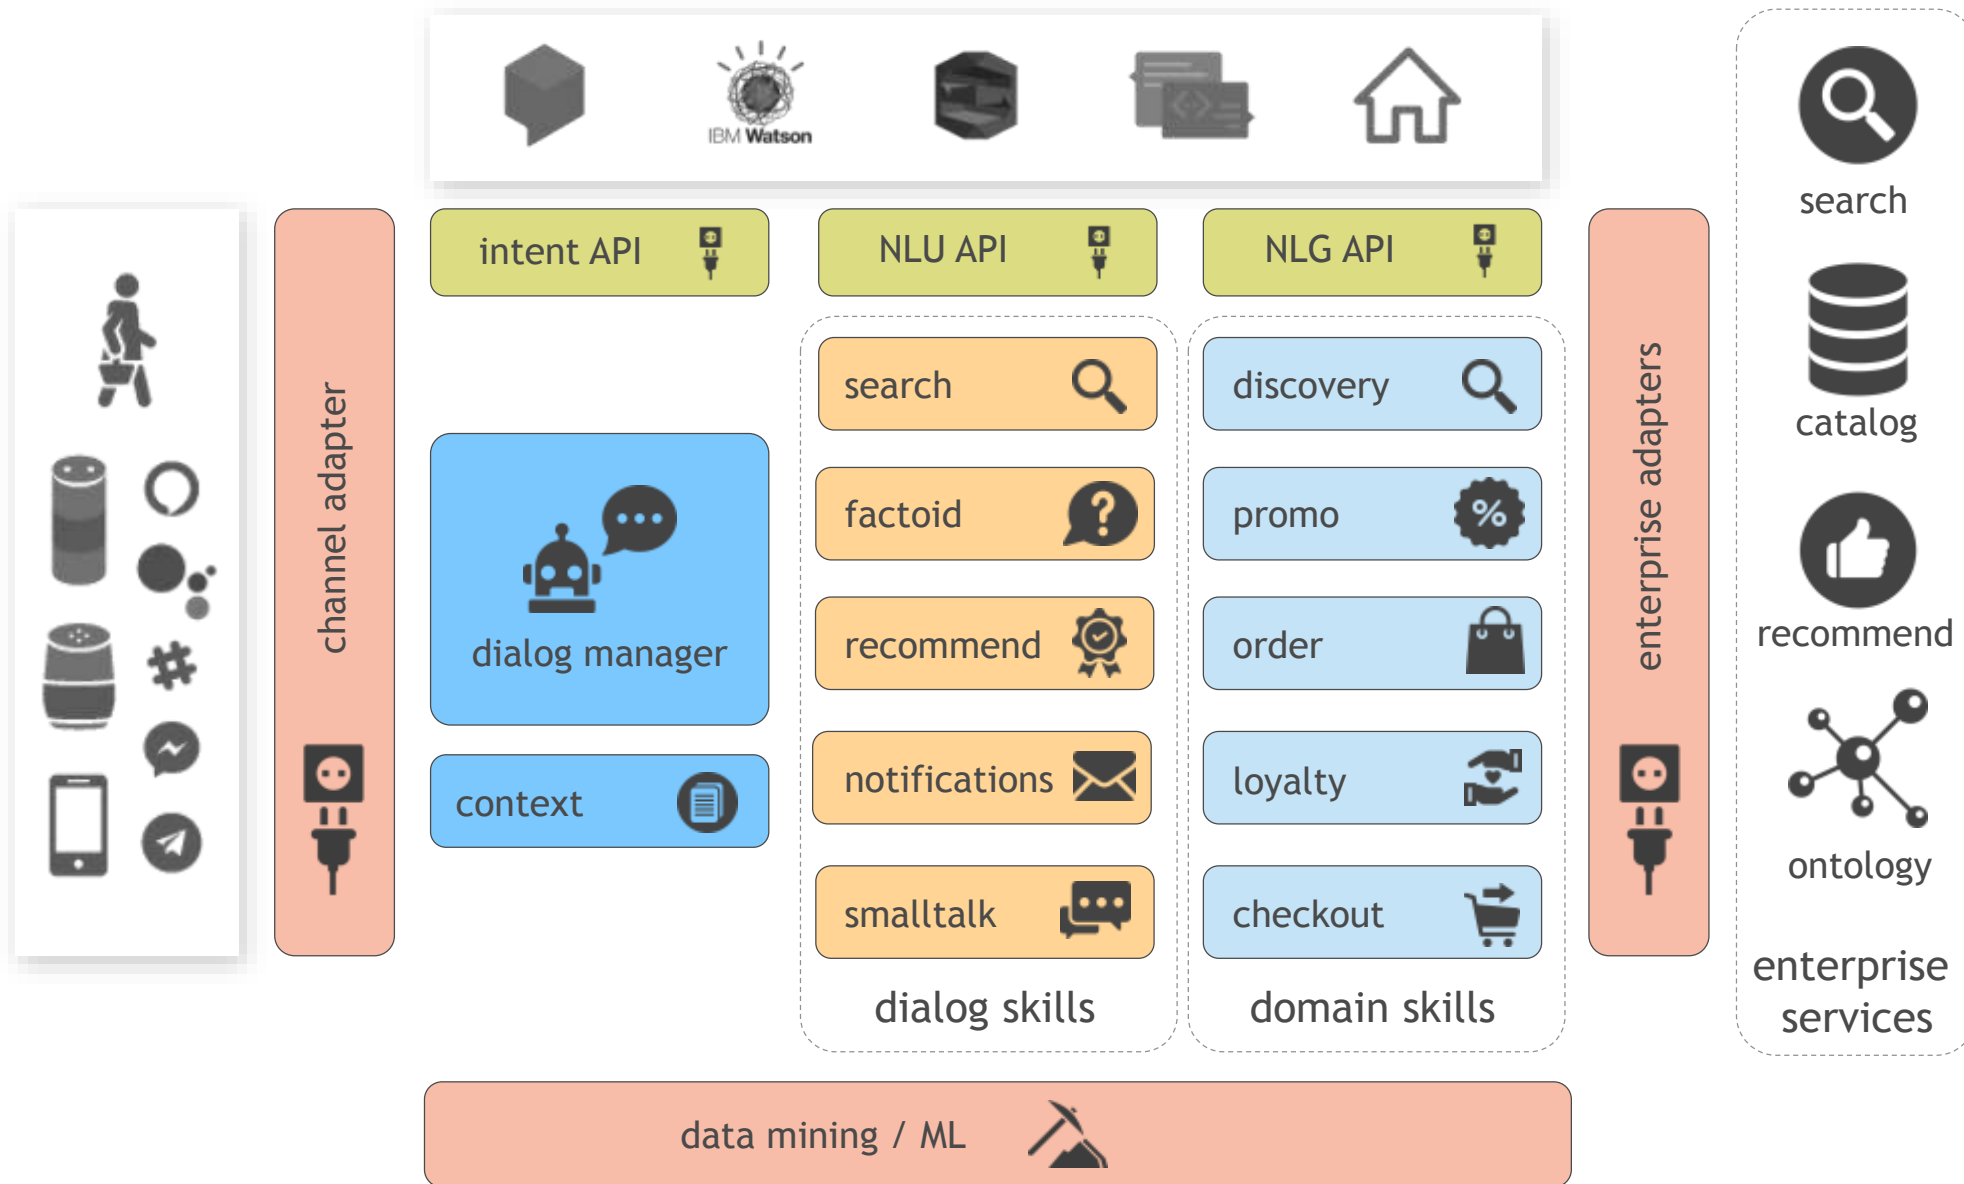
Why Grid Genie project?

- Demonstrate technical feasibility of current state-of-the-art
- Explore the limits of current technologies
- Learn to develop, test, deploy CUI applications
- Develop reference designs and tooling for all facets of conversational application development
- Accelerate development of new CUI applications

Conversational demo design principles

- **100% open and free:** written using open source technologies in pluggable architecture
- **Multiple devices:** supports Google Home, Alexa, Siri, etc. via device adapters
- **Multiple channels:** voice-based, text-based, web-based via channel adapters
- **Multiple skills:** specific dialog skills can be added (promo, checkout, order tracking)
- **Multiple models:** mix, match and use best-of-breed ML/NLU models for specific tasks
- **Integrated with backend services** via layers of platform services and enterprise adapters
- **Containerized:** for seamless on premise or cloud deployment

Conversational commerce blueprint



Let's chat about cameras

Why cameras?

- Diverse category of products: about 1,000 items in catalog
- Wide range of subcategories, from cheap point-n-shoot to high-end professional cameras
- Technical products with vastly different features
- Some folks know exactly what they want; others don't know much about them

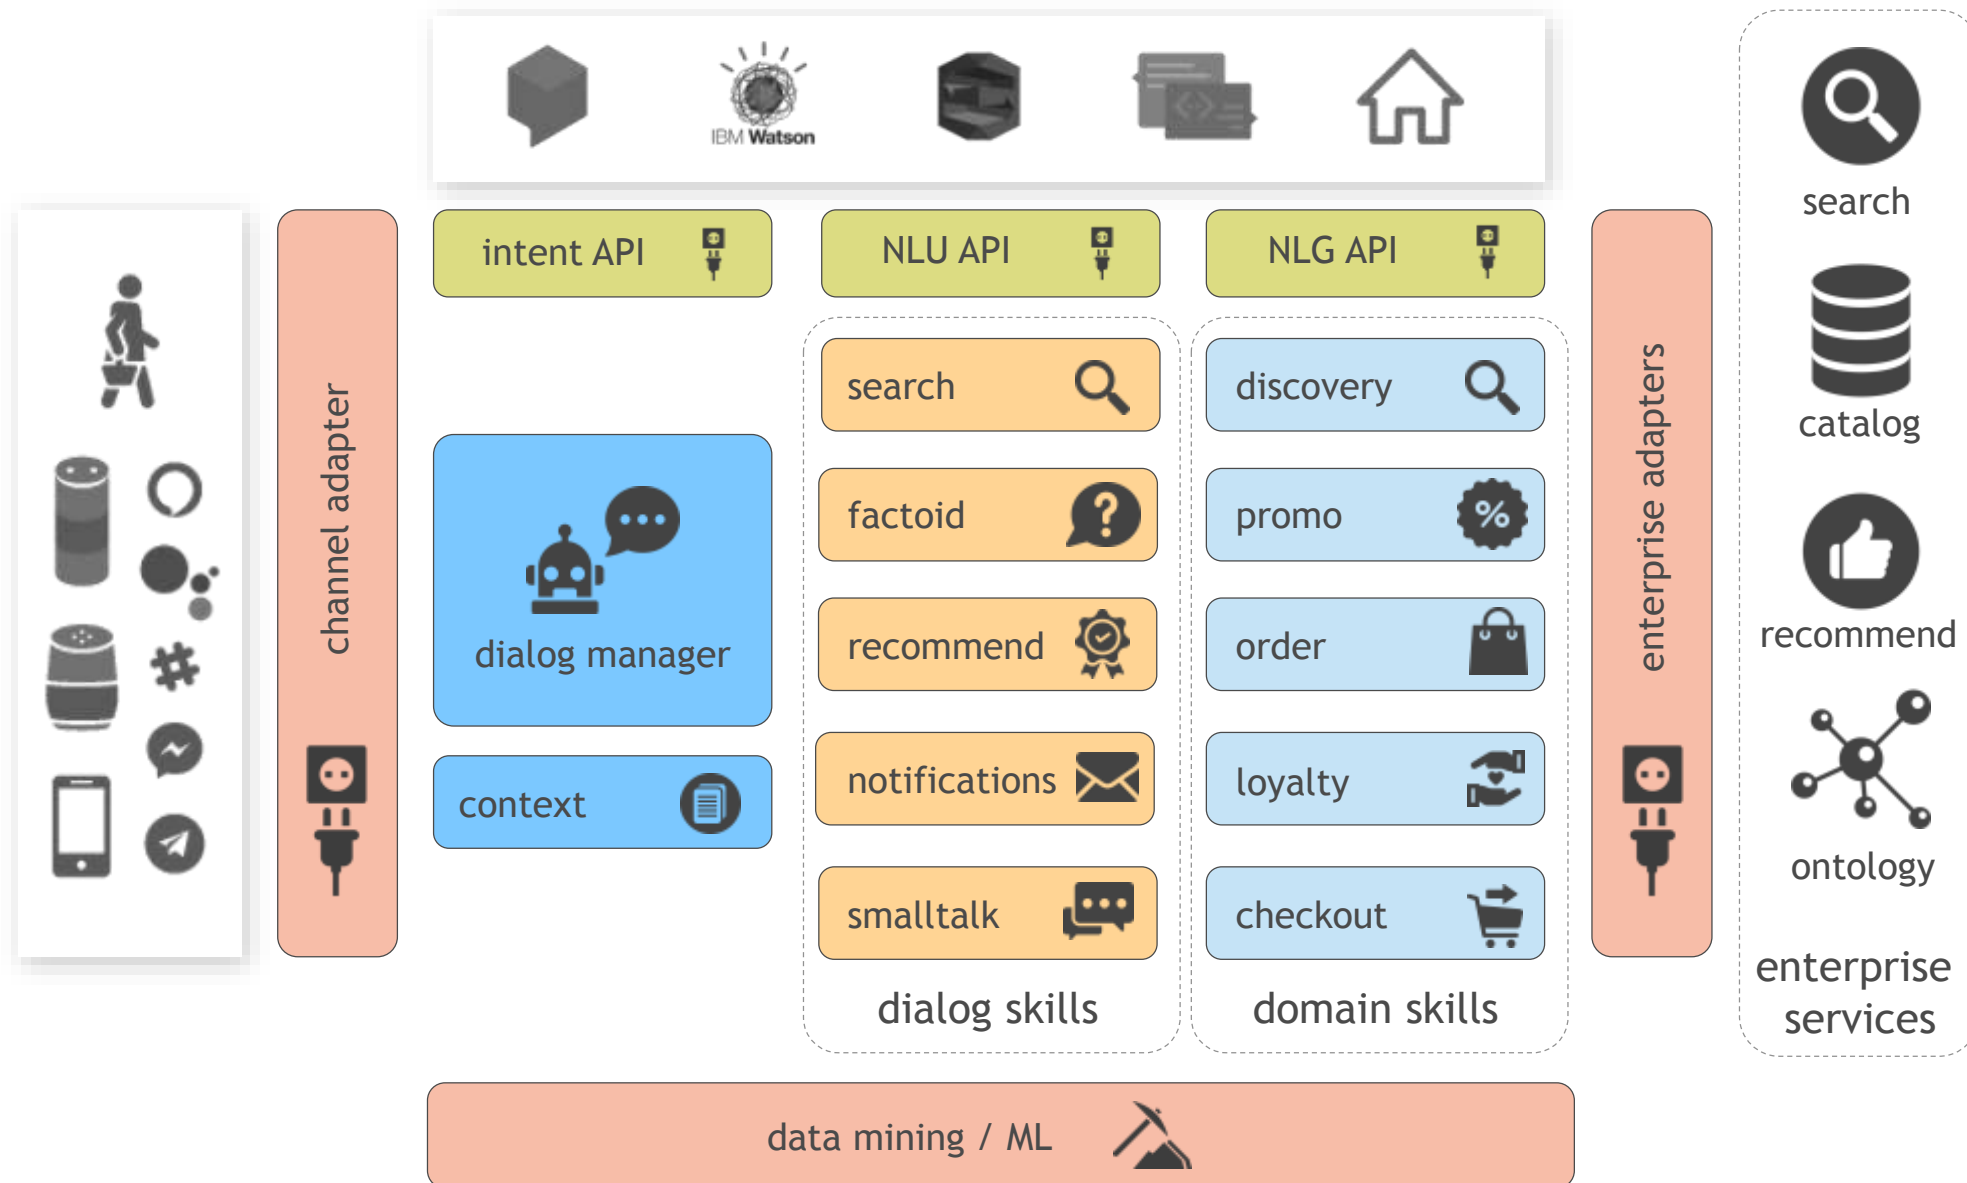
What we are looking to proof?

- Device interoperability
- One dialog manager arbitrating between multiple (pluggable) skills
- Taxonomy-based domain knowledge
Better dialog thru better taxonomy
- Know the difference between advisory & order taking
- Fluidity of conversation, keeping it natural, human-like

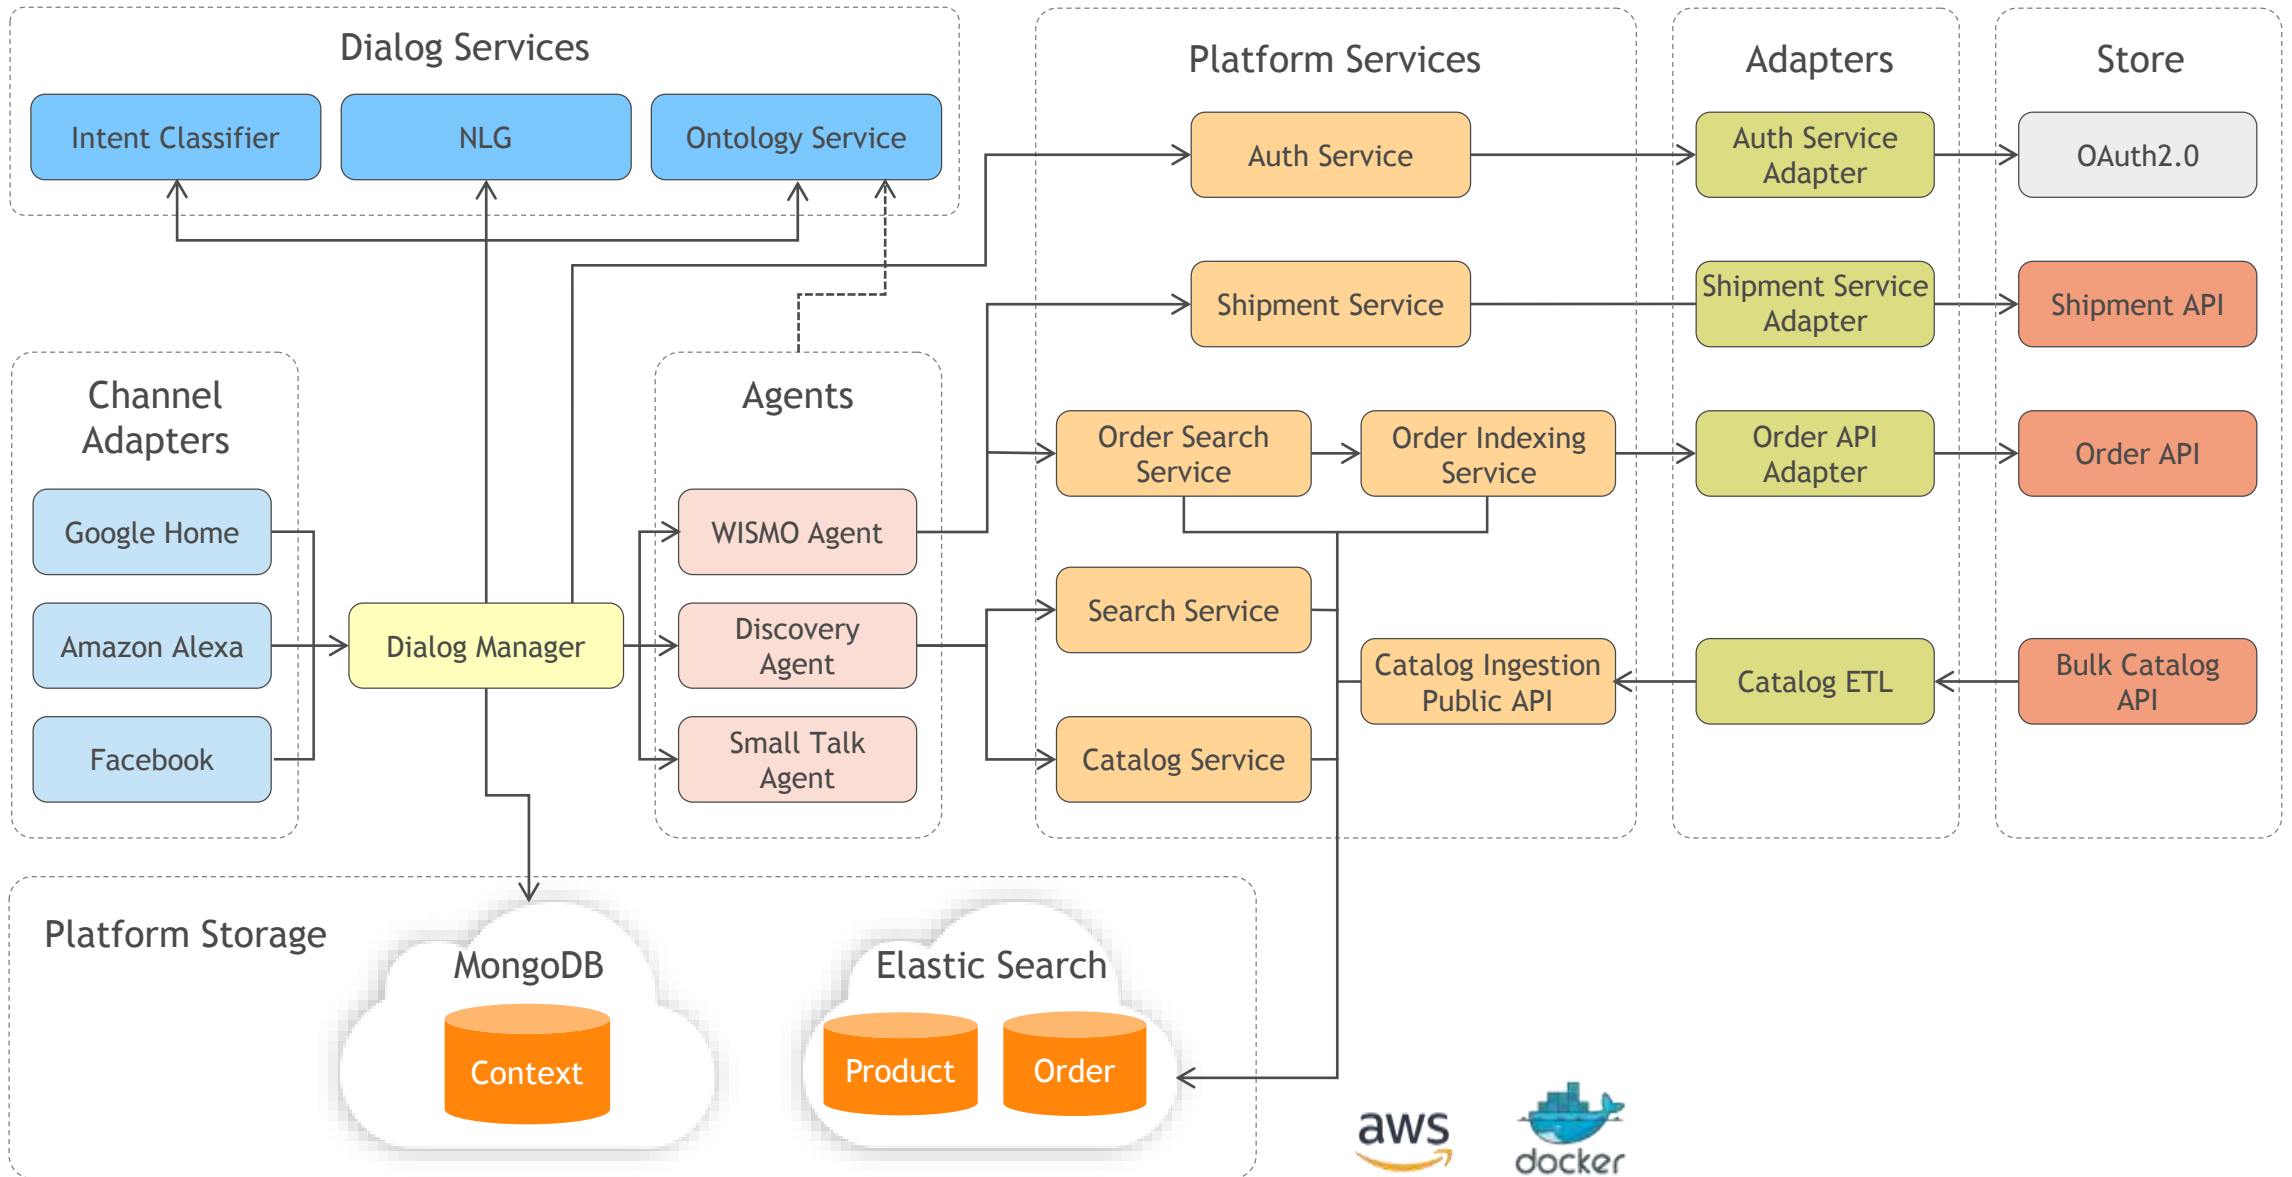
Grid Genie Demo



Conversational commerce blueprint



Microservices architecture



Lessons learned (so far)?

- Reference architecture works (by and large)
- Closed platforms like Google dialog flow, are easy to start with, but hard to grow with
- It is possible to write open, cross-device applications!
- But... lack of standards leave you at mercy of (frequently) changing Google/Alexa APIs
- So, continuous testing (and full true CICD) is key to spot regressions quickly
- Google and Alexa publishing standards help test your application
- Creating pluggable taxonomies is hard
- Finding good dialog writers is hard
- Best practices for writing general-purpose, cross-platform dialog managers are immature

- Conversational applications are ready for prime time
- It's fun and practical to write them, although not (yet) cheap or easy

Q&A




Conversational platform components

Service	Description	Responsibilities
Channel adapters	Connect the platform to edge devices	Account linking, Auth, SSO, user management
Context	Context keeps dialog state	Shares context between devices and interactions
Dialog manager	Controls conversation flow	maintains conversation context, detects the customer's broader intent, delegates tasks to an appropriate dialog agent
Pluggable dialog agents	Implement narrow conversation flow	execute particular use case, communicate with NLU, enterprise adapters to fulfill the actions, use NLG to form natural language response
Dialog services	NLU/NLG/NER service adapters and implementations	Provide NLG/NLU/NER capabilities
Enterprise adapters	Connect agents to data and services	Provide access to enterprise APIs
Domain ontology	Capture domain knowledge	deep understanding of business domain, its terminology, slang, and the relationship between terms

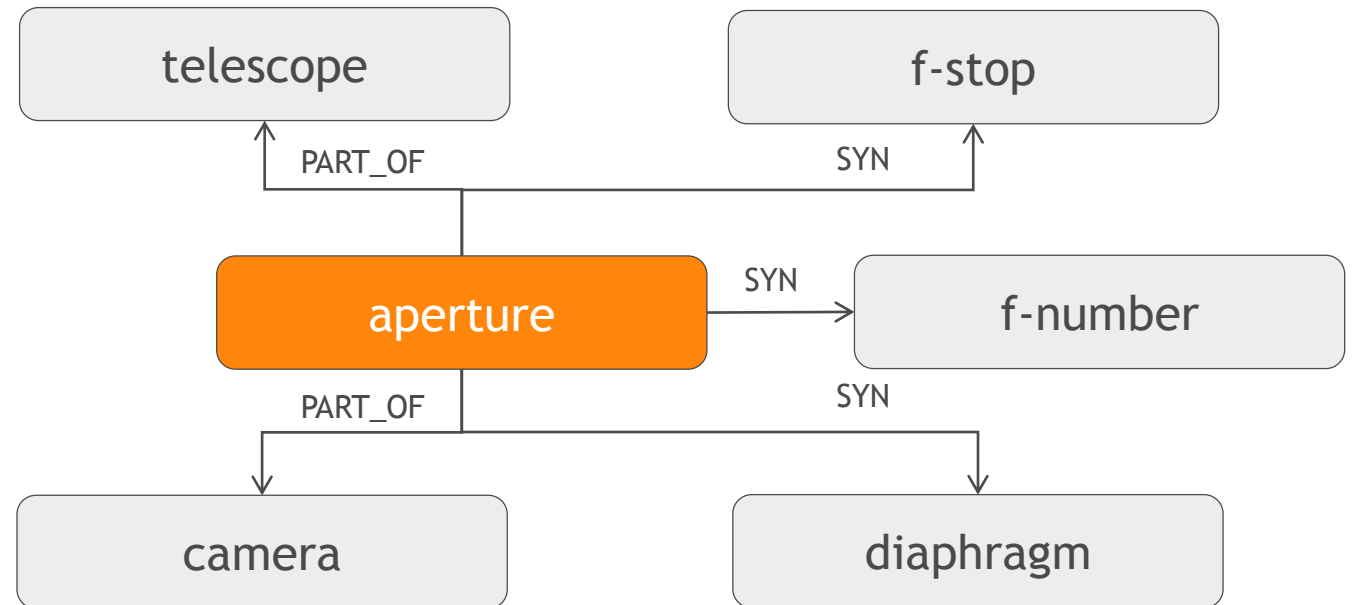
Domain specific ontology

Implementation:

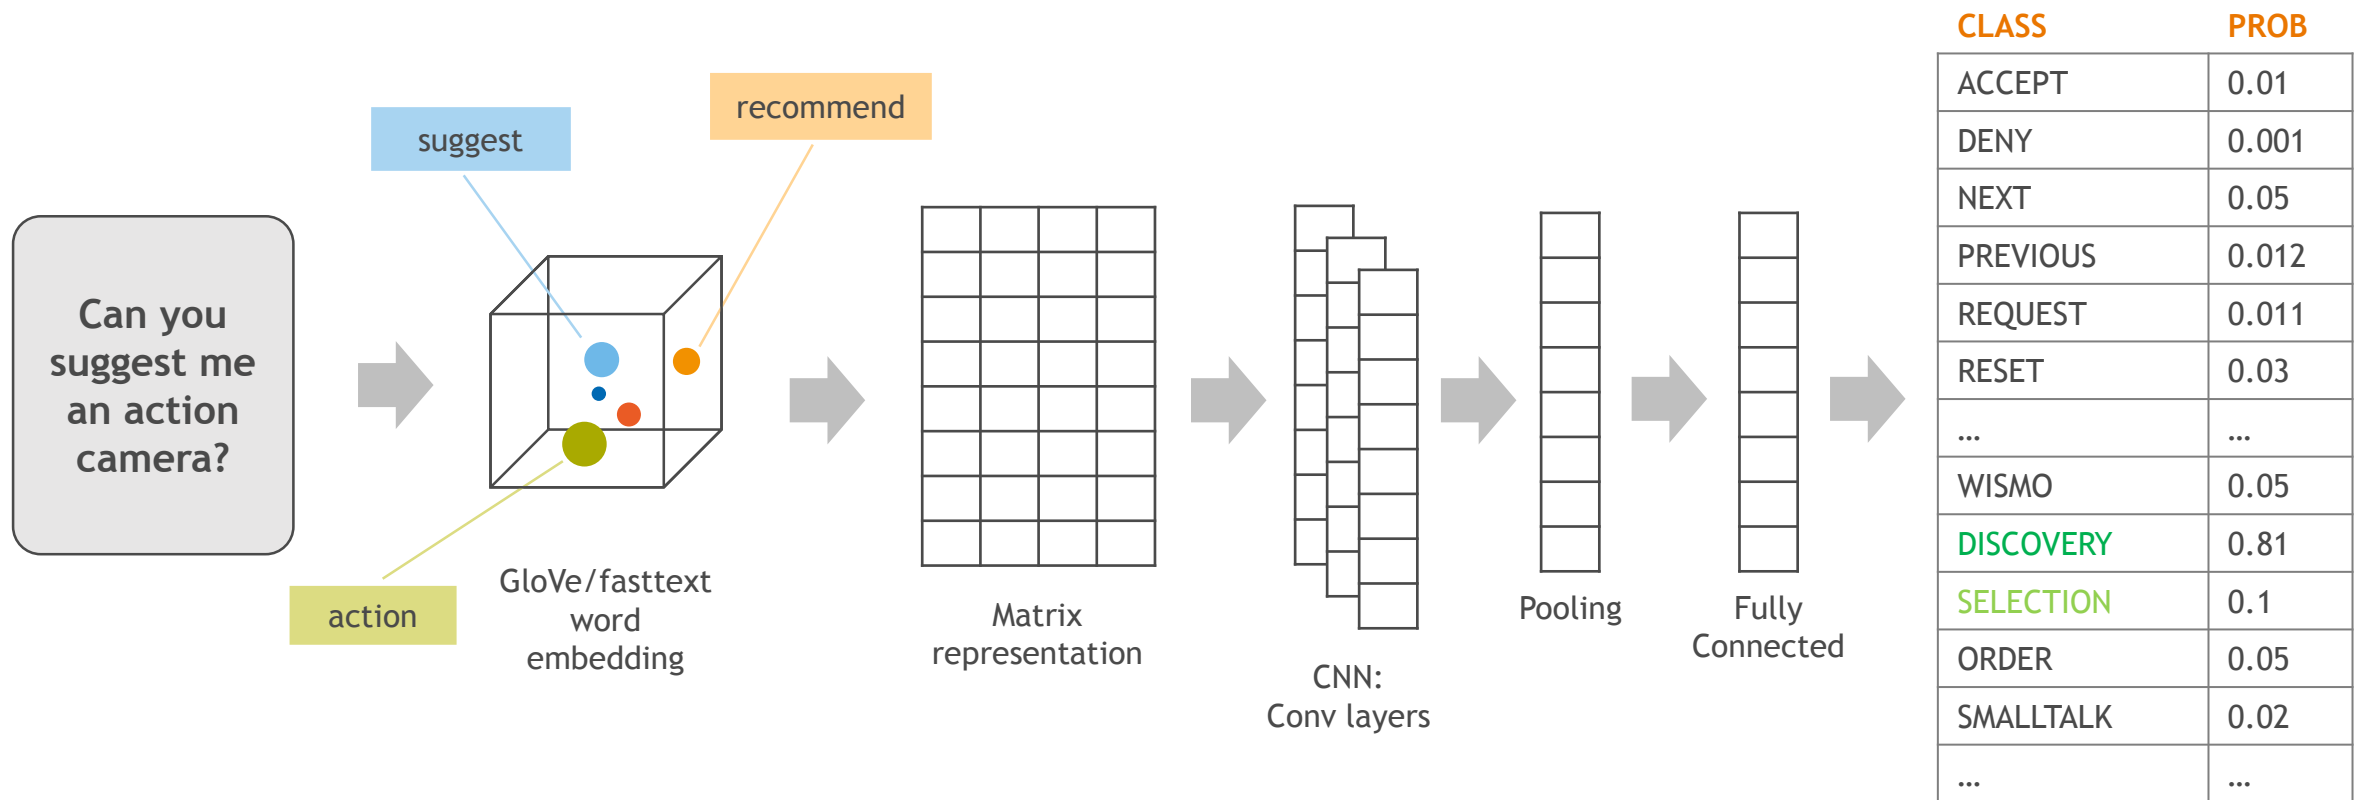
- Neo4J 

Content sources:

- BabelNet 
- WordNet 
- DBpedia 



Intent classification



CNN based model for intent classification

To learn more

- Subscribe to our blog: blog.griddynamics.com

Appendix

1 1 0 1 1 1 0 1 1 1 1
0 1 0 1 1 1 1 0 0 0 0 1 1 1 1 1 1 0 1
1 0 1 0 1 0 1 1 0 1 1 1 1 0 1 0 1 1 1 1 0 1
1 1 0 1 0 0 0 0 1 0 1 1 0 1 1 0 1 1 1 1 1 0
1 0 1 0 1 0 1 1 0 1 0 1 1 1 1 0 0 1 0
1 0 1 0 1 0 1 0 1 1 1 1 1 1 0 1 0 1 1 0 0 1 0 1
0 1 0 1 0 1 1 1 1 0 1 0 1 1 0 0 1 0 1 1 0 1 1 1 1 0
0 1 1 1 1 1 0 1 0 1




0 1 1 1 0 1 0 1 0 1 1
0 1 0 1 0 1 1 1 1 0
0 0 1 0 1 0 1 1
0 1 0 1 0 0 1 0 1 0
0 1 0 1 0 1 0 1 0 1 0

Named entity recognition: BIO tagging

I	need	diving	camera	under	two	hundred	bucks
O	O	B-purpose	B-device_name	B-relational	B-quantity	I-quantity	B-unit

BIO tags (Beginning, Inside, Outside)



BiLSTM based model for entity tagging

Named entity recognition: tag grouping

I	need	diving	camera	under	two	hundred	bucks
O	O	B-purpose	B-device_name	B-relational	B-quantity	I-quantity	B-unit



Linear XGBoost model is trained to make binary classification for token pairs

About Grid Dynamics



Founded in 2006, Grid Dynamics is an engineering services company built on the premise that cloud computing is disruptive within the enterprise technology landscape. Since that time, we've had the privilege to help companies like Microsoft, eBay, PayPal, Cisco, Macy's, Yahoo, ING, Bank of America, Kohl's, among others, to re-architect their core mission-critical systems, develop new cloud services, accelerate innovation cycles, increase software quality, and automate application management.

Grid Dynamics has multiple locations in the USA and Europe, and employs over 1,000 expert engineers worldwide.



Thank you!



www.griddynamics.com