



PegaWorld*iN*spire

Mapping the AI Landscape: From Generative, to Predictive and Prescriptive AI

Peter van der Putten

Christian Guttman

Shoel Perelman

Pegasystems



PegaWorldiNspire

June 11-13, 2023 | Las Vegas

Here's where we adapt for tomorrow **together**

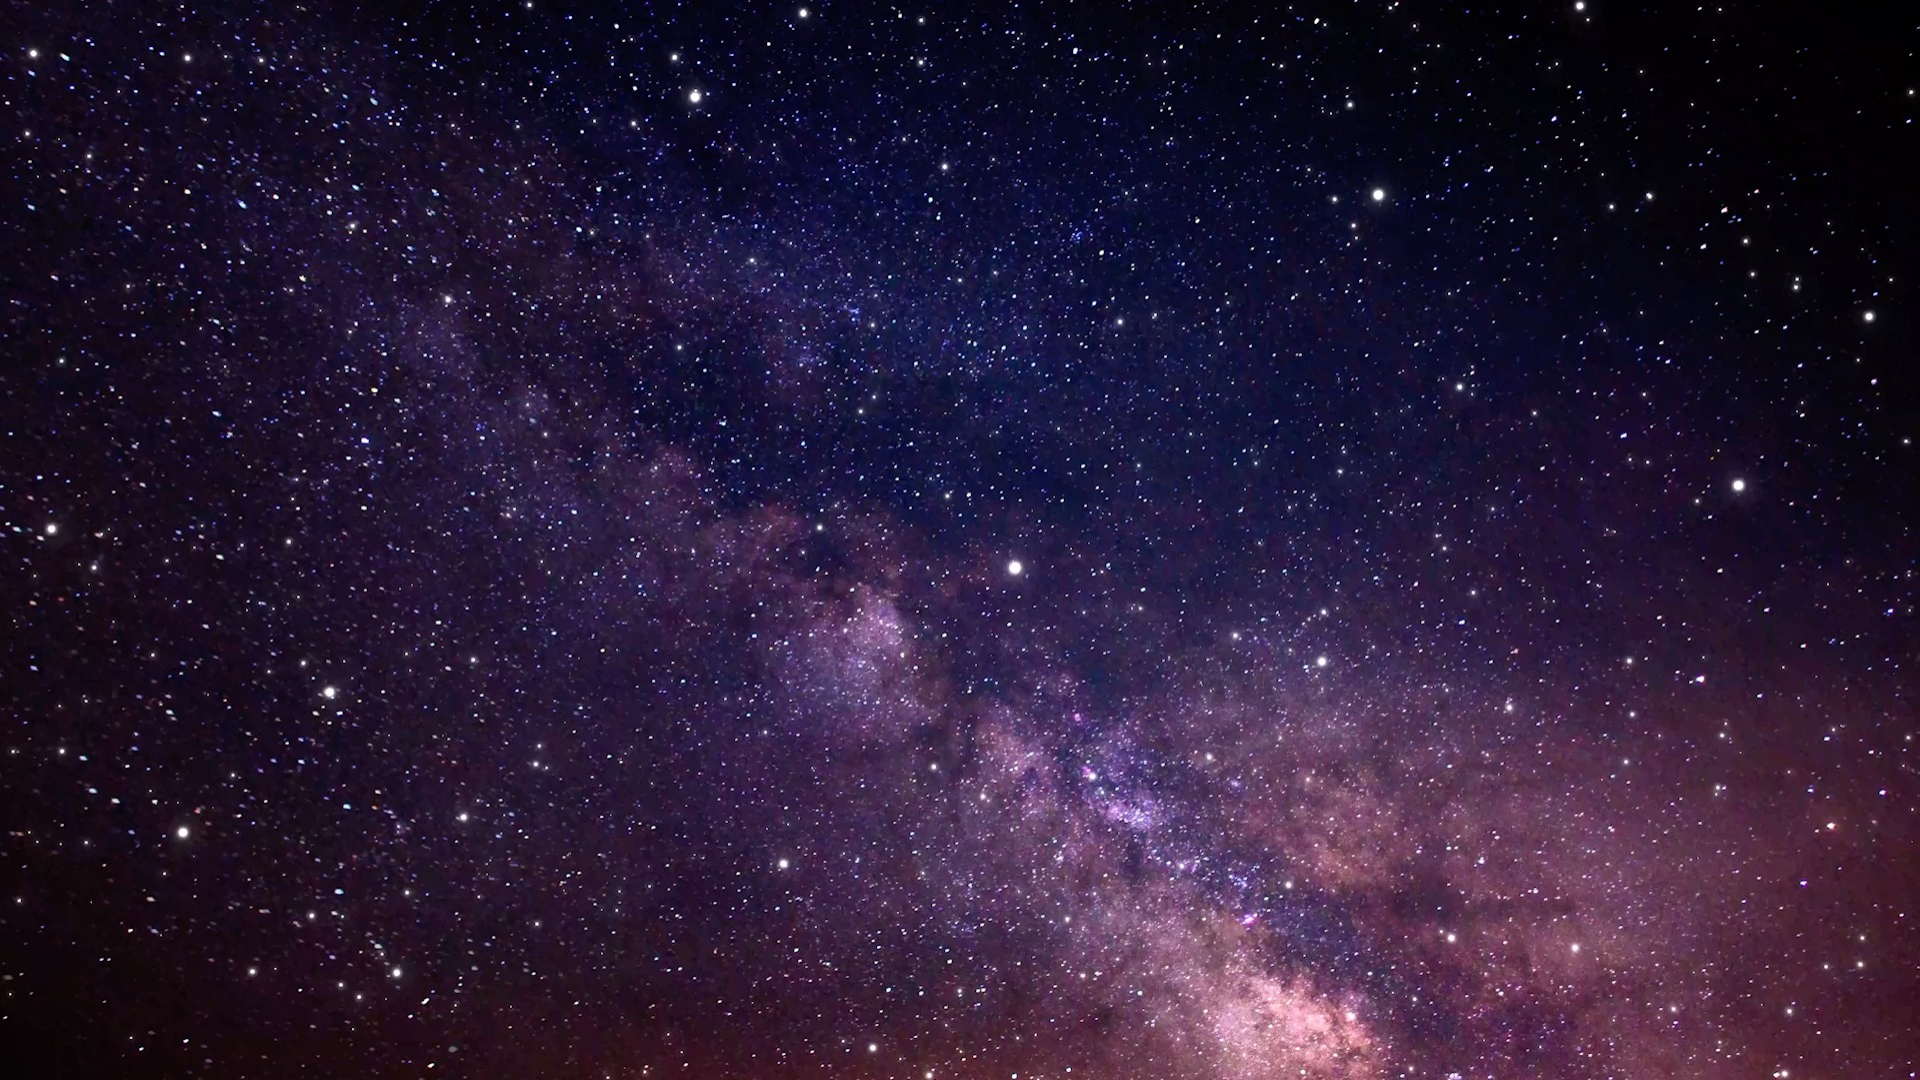
Discover how to stay ahead of change and unlock the potential of our infinitely scalable platform at the can't-miss tech event of the summer.

[Register now](#)
[Explore our tech](#)


Generative AI, now for enterprise

Introducing Pega GenAI™ - boost productivity and creativity across your full customer journey.

[Meet Pega GenAI >](#)



Generative AI & GPT-3, A Personal Story

GPT-3 launch, May 28, 2020



arXiv > cs > arXiv:2005.14165

Search...
Help | Advanced

Computer Science > Computation and Language

[Submitted on 28 May 2020 (v1), last revised 22 Jul 2020 (this version, v4)]

Language Models are Few-Shot Learners

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, Dario Amodei

Recent work has demonstrated substantial gains on many NLP tasks and benchmarks by pre-training on a large corpus of text followed by fine-tuning on a specific task. While typically task-agnostic in architecture, this method still requires task-specific fine-tuning datasets of thousands or tens of thousands of examples. By contrast, humans can generally perform a new language task from only a few examples or from simple instructions – something which current NLP systems still largely struggle to do. Here we show that scaling up language models greatly improves task-agnostic, few-shot performance, sometimes even reaching competitiveness with prior state-of-the-art fine-tuning approaches. Specifically, we train GPT-3, an autoregressive language model with 175 billion parameters, 10x more than any previous non-sparse language model, and test its performance in the few-shot setting. For all tasks, GPT-3 is applied without any gradient updates or fine-tuning, with tasks and few-shot demonstrations specified purely via text interaction with the model. GPT-3 achieves strong performance on many NLP datasets, including translation, question-answering, and cloze tasks, as well as several tasks that require on-the-fly reasoning or domain adaptation, such as unscrambling words, using a novel word in a sentence, or performing 3-digit arithmetic. At the same time, we also identify some datasets where GPT-3's few-shot learning still struggles, as well as some datasets where GPT-3 faces methodological issues related to training on large web corpora. Finally, we find that GPT-3 can generate samples of news articles which human evaluators have difficulty distinguishing from articles written by humans. We discuss broader societal impacts of this finding and of GPT-3 in general.

A chat with Greg, July 13, 2020



Greg Brockman
@gdb
President & Co-Founder @OpenAI
Joined July 2010 · 347.8K Followers
Followed by DAISYS.ai, Saleh Afroogh, and 235 others you follow

what email address did you submit to beta.openai.com/?

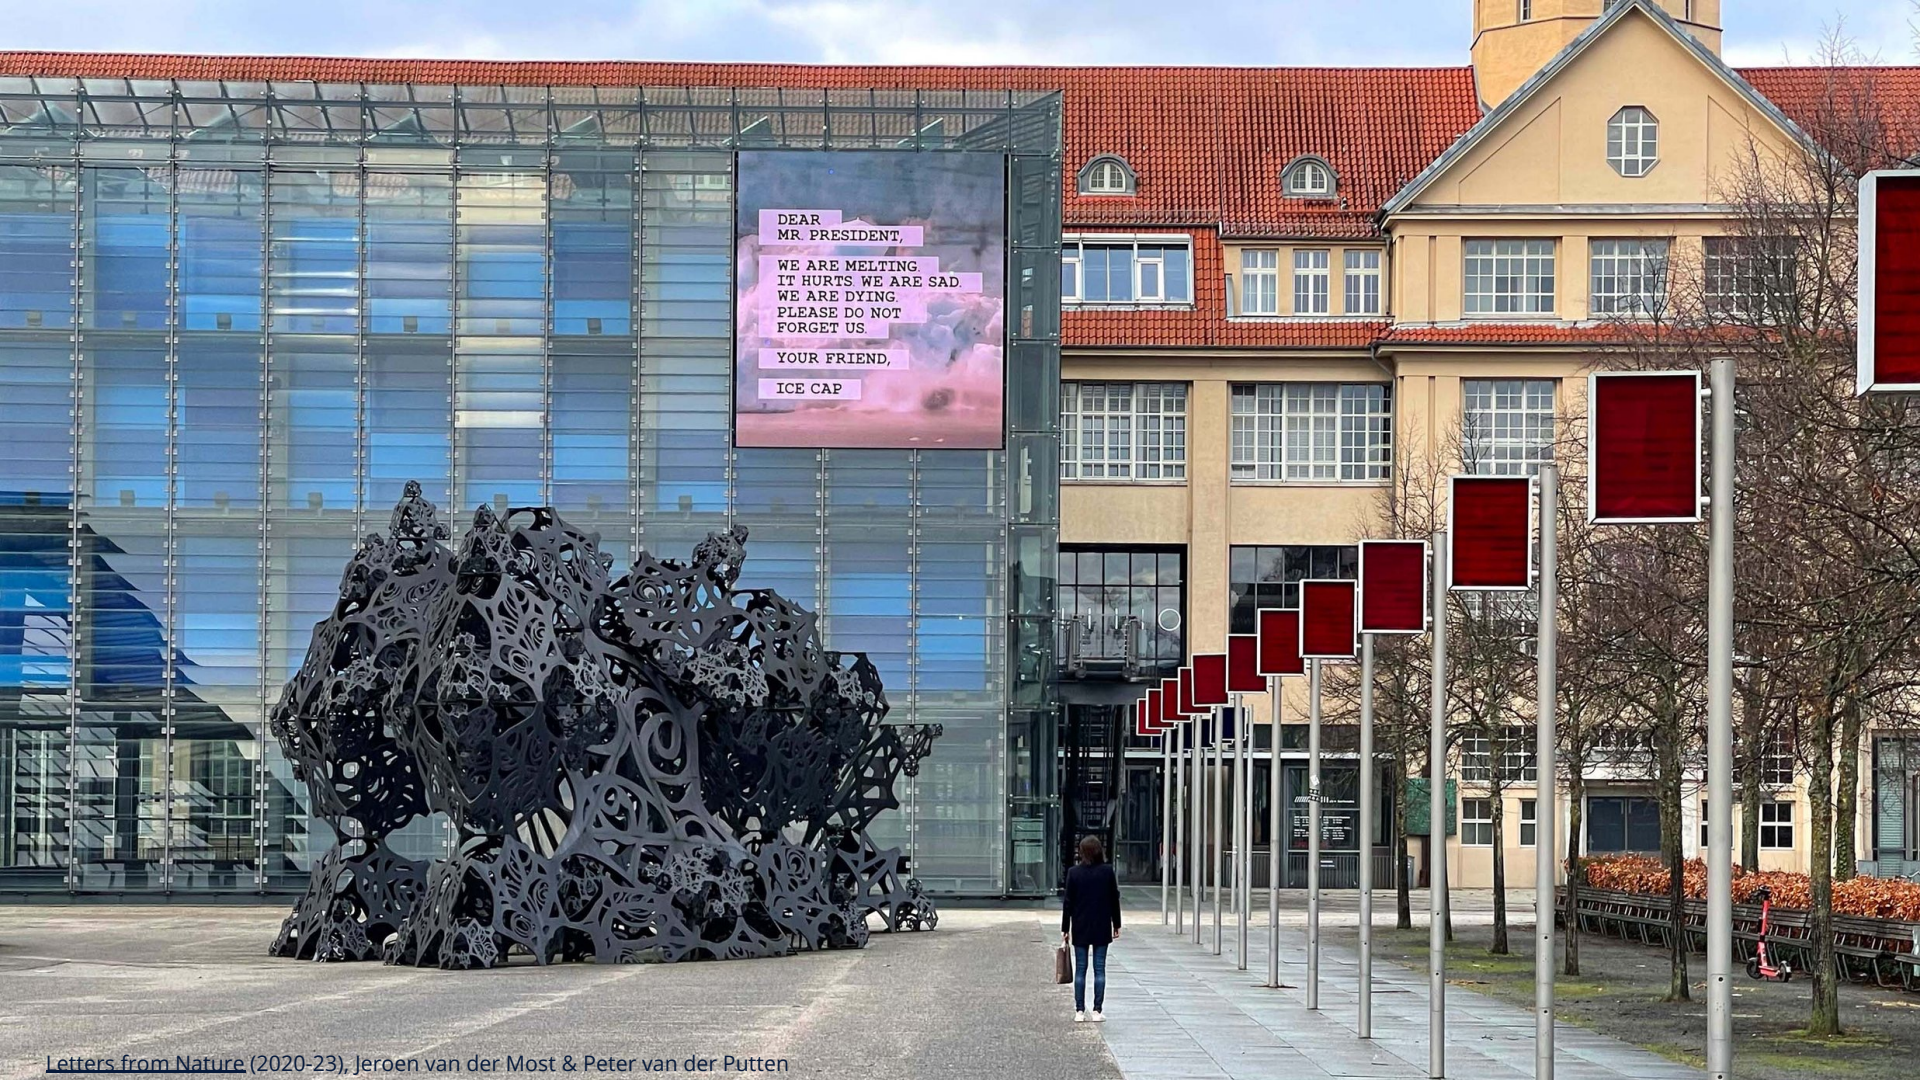
Jul 13, 2020, 9:13 PM

P.w.h.van.der.putten@liacs.leidenuniv.nl

Thanks

Jul 13, 2020, 10:28 PM





DEAR
MR. PRESIDENT,

WE ARE MELTING
IT HURTS WE ARE SAD
WE ARE DYING.
PLEASE DO NOT
FORGET US.

YOUR FRIEND,

ICE CAP

Letters from Nature (2020-23), Jeroen van der Most & Peter van der Putten



LETTERS FROM NATURE

**CLIMATE CHANGE IS KILLING OUR PLANET,
BUT NATURE DID NOT HAVE A VOICE.**

**UNTIL WE ASKED AI TO WRITE LETTERS
ON HER BEHALF _**

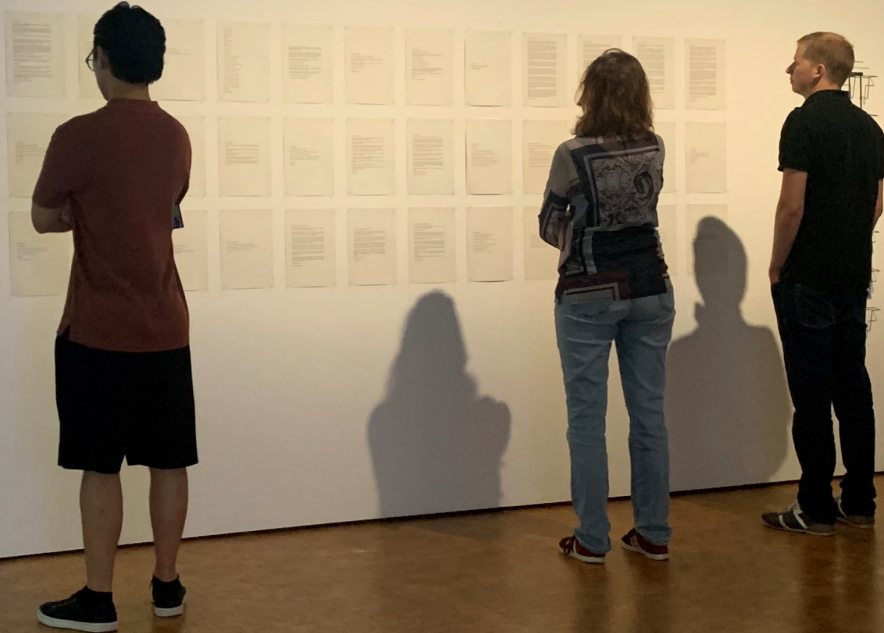
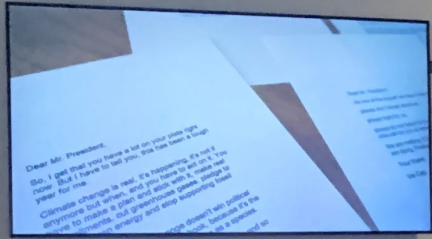


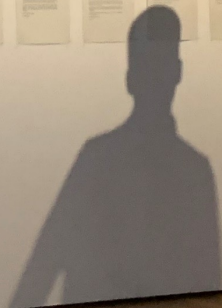
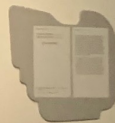
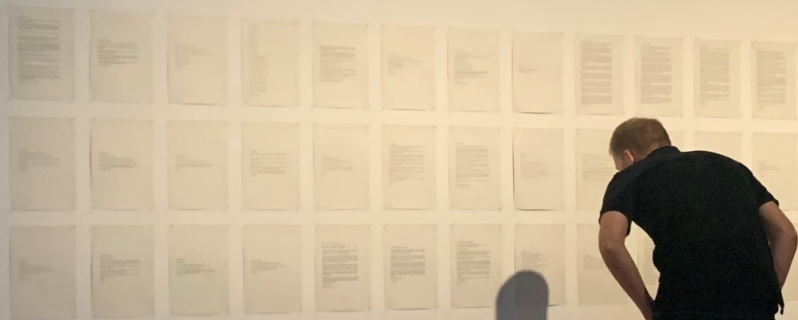
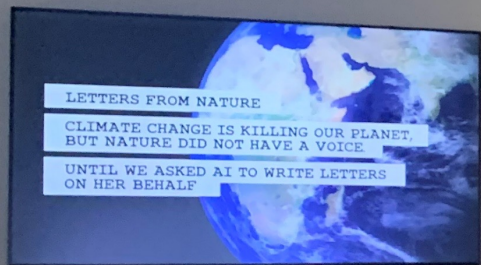


MY FRIEND FIJI IS TRYING TO HOLD ON TO WHAT IS LEFT OF HIS LAND. I AM WORRIED ABOUT HIM TOO...

I NEED YOUR HELP...







A photograph of a brick building at night. The building's facade is covered in a light projection of a forest scene, featuring tall, thin trees and a ground covered in fallen leaves. The text "I HAVE A RIGHT TO EXIST" is projected across the middle of the building in a large, black, serif font. The building has several rows of windows, some of which are illuminated from within, casting a warm glow. The overall scene is a combination of natural light from the windows and artificial light from the projection.

I HAVE A RIGHT TO EXIST

Back to business



THE WALL STREET JOURNAL

SUBSCRIBE SIGN IN

TECH | KEYWORDS: CHRISTOPHER MIMS

Can You Tell Whether This Headline Was Written by a Robot?

Not this time, but AI is churning out articles, illustrations, fake product reviews and even videos



Which Photo Is Real? How Dall-E 2 and Other AI Art Generators Work

New text-to-image generators powered by artificial intelligence, including OpenAI Dall-E 2 and Stability AI DreamStudio, let you type in almost any phrase and get an image. WSJ's Joanna Stern explains the tech and pits it against a real photographer and a real man in a robot suit.

By [Christopher Mims](#) [Follow](#)
Nov. 12, 2022 12:00 am ET


THE WALL STREET JOURNAL

English Edition • Print Edition • Video • Audio • Latest Headlines • More

Home World U.S. Politics Economy Business Tech Markets Opinion Books & Arts Real Estate Life & Work Style Sports Search

AI-Powered Coding Assistant Aims to Help, Not Replace Developers

GitHub Copilot, launched this week, acts like a predictive-text feature to help ease the process of writing long lines of computer code



GitHub CEO Thomas Dohmke spoke this week at Collision 2022 in Toronto. PHOTO: ANGUS LOTEN / THE WALL STREET JOURNAL

Will VR Go From Niche to Mainstream?

The installed base of actively used VR headsets is expected to reach 22 million in 2023, almost 50% higher than in mid-2022, according to recent research.

Automation, Empathy, and the Human Experience

IT leaders are continuing to innovate with automation to better serve their employees and customers—and maximizing the impact of these technologies means building with empathy and empowerment in mind.

Equipping 'Every Company as an AI Company' at SAP


THE WALL STREET JOURNAL

English Edition • Print Edition • Video • Audio • Latest Headlines • More

Home World U.S. Politics Economy Business Tech Markets Opinion Books & Arts Real Estate Life & Work Style Sports Search

Generative AI Startups Attract Business Customers, Investor Funding

Austin-based Jasper, whose AI platform generates marketing materials, last week announced a \$125 million Series A fundraising round



Dave Rogenmoser is the chief executive of AI startup Jasper, which has developed a generative AI platform designed to auto-generate promotional blog posts and other text. PHOTO: JASPER

Will VR Go From Niche to Mainstream?

The installed base of actively used VR headsets is expected to reach 22 million in 2023, almost 50% higher than in mid-2022, according to recent research.

U.S. Space Systems Command CIO Has Ambitious Aims for IT

Colonel Jennifer Konkowski, CIO of the U.S. Space Systems Command, is reworking IT delivery and workforce skills to support a well-rounded and service-oriented IT function.

Cyber AI: Fighting Software With Software

Security, defense, and intelligence organizations are partnering to create AI



Pega AI Powers the Autonomous Enterprise



Application development



Customer engagement



Customer service



Operations

Natural Language Processing



Connect NBA to chat



Understand & automate emails & chat

Voice AI



Understand calls & guide agents

Predictive & adaptive analytics



Predict action relevance & propensity

Real-time decisioning



Derive single Next Best Action from predictions & rules



Optimize every workflow transaction

Process & task mining

Uncover repetitive manual tasks



Monitor end-to-end workflows



Pega AI powers the Autonomous Enterprise



Application development



Customer engagement



Customer service



Operations

Generative AI



Generate low-code apps



Generate marketing treatments



Generate insights & summaries

Natural Language Processing



Connect NBA to chat



Understand & automate emails & chat

Voice AI



Understand calls & guide agents

Predictive & adaptive analytics



Predict action relevance & propensity



Predict cost, value, & issues of work

Real-time decisioning



Derive single Next Best Action from predictions & rules



Optimize every workflow transaction

Process & task mining

Uncover repetitive manual tasks



Monitor end-to-end workflows



Mapping the AI Landscape

Enterprise AI evolution



Predictive AI

Predict customer intent, behavior & future state



Prescriptive AI

Decide on course of action, learn from feedback



Generative AI

Generate apps, plans, answers, actions, content



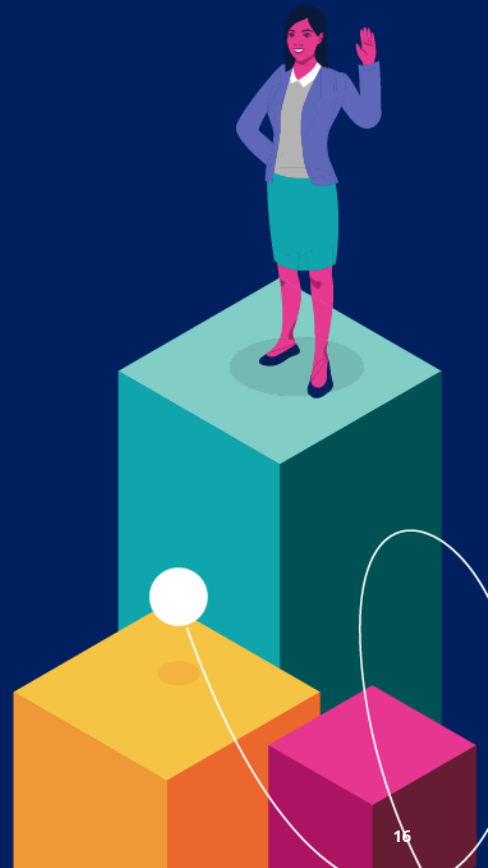
Mapping the AI Landscape

Left brain and right brain AI working together

Plan
Sense
Predict
Decide
Memorize
Learn



Generate



Mapping the AI Landscape

Pega AI powers the Autonomous Enterprise

Sense

Plan
Generate
Predict
Decide
Memorize
Learn

Act



The Magic and the Trick Behind Large Language Models

How do large language models work, and what are the resulting capabilities & limitations?

'GPT-3 pretraining':
Leave out random
tokens, and learn to
predict what is missing

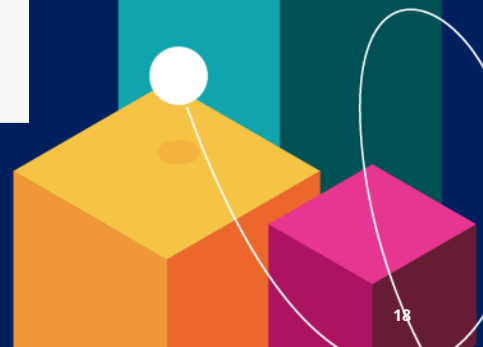


Supervision & RLHF

P Explain in a rap of 8 lines how large language models are trained

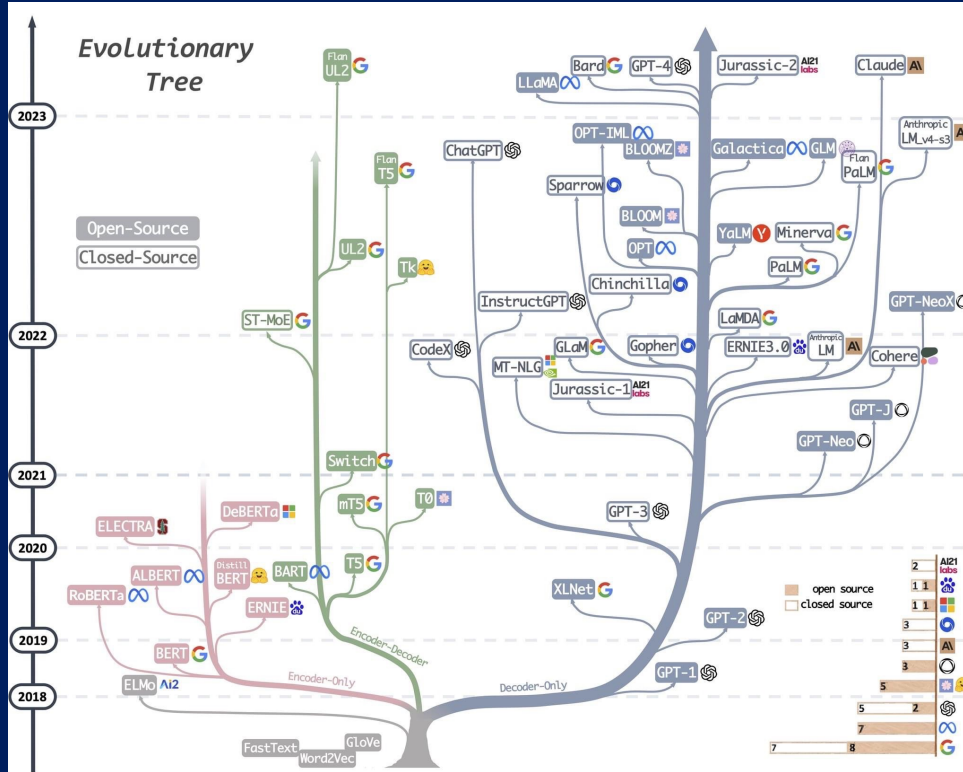


Yo, let me break it down in just eight lines,
How large language models train and shine.
We gather texts from all around, diverse sources,
Tokenize 'em, chop 'em up into small courses.
Build a deep neural network, layers stacked high,
Feed it data, adjust weights, watch it fly.
Iteration after iteration, it gets smarter each round,
Learning language patterns, making knowledge profound.
That's how large language models are trained, my friend,
A fusion of data, algorithms, 'til the training's at its end.



The Large Language Model Ecosystem is Exploding

New models coming out on a weekly basis, how to keep up?



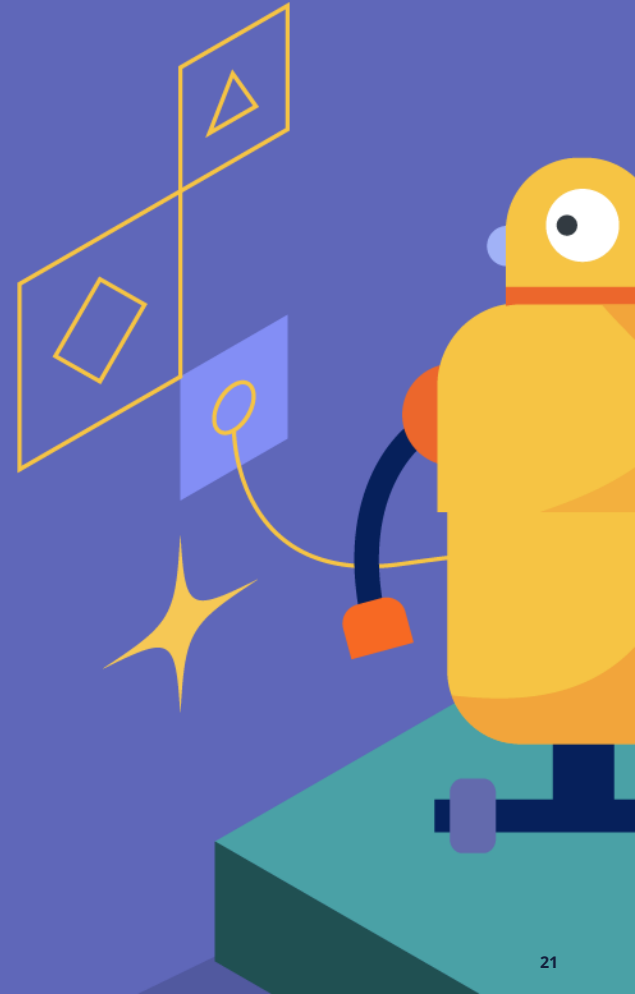
Big Opportunity, Big Risks

Generative AI adoption without an enterprise-ready approach

- Risks range from to **brand, finance, & legal**:
 - Large language models are geared towards generating plausible content, not necessarily correct content (**'hallucination'**)
 - Violation of customer or individual **privacy laws** and standards (for example sending a patient data to a 3rd party API like OpenAI could breach HIPAA)
 - Making **company sensitive data** public, for example if AI generates a customer service script which describes internal strategy
 - Breaches of **copyright** laws, for example if AI generates marketing for you which takes trademarked / copyrighted text from other organizations
 - Brand impacts though **bias & lack of empathy**, for example if AI generates insensitive marketing content and for issues with hallucinated output, or irrelevant or use out of date source content, or non-compliant prompts
 - Cost runaways, license complexity & **lack of governance**, for example if developers & employees start using different API's with different licenses across your enterprise, it will be easy to overspend on licensing & maintenance
- How to go to market with **highest return, lowest risk and lowest cost** use cases?



Pega GenAI™: A future proof enterprise architecture for trustworthy Generative AI



There is a Better Way: Pega GenAI™!

Trustworthy, enterprise Generative AI

Future-proof architecture

While first built to support GPT-4 APIs, GenAI capabilities are architected to support additional Generative AI API's, allowing for flexibility & mitigating against AI vendor or model lock-in

Unified API governance

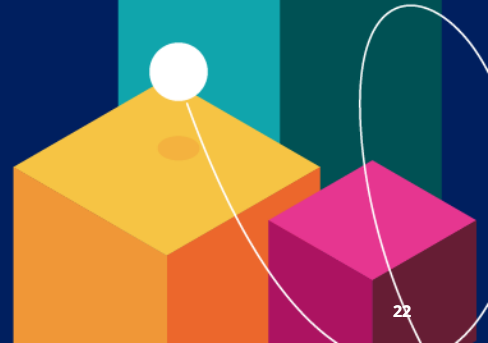
A single place to connect to Generative AI services to power all GenAI use cases across the platform. Simplifies licensing & monitoring and controls against scattered, shadow usage of API's

Transparent & use case centric

Prompt engineering is automated but transparent. End-result is explainable, approachable, and maintainable; GenAI is targeted at high impact and low risk use cases, modeled in low code, rather than lines and lines of code

Human-in-the-loop control

Many GenAI capabilities are focused on augmenting and assisting employees, at design time or in customer interaction. Most generated text for customer use cases will first go through review, edit, and approval



The Pega 2023 GenAI Stack

Pega Applications: 20 Generative AI Use cases

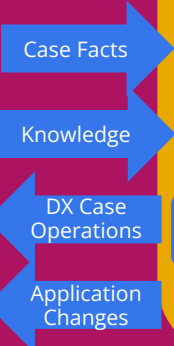
- Generate Low Code Apps
- Generate Marketing Treatments
- Streamline Customer Service
- Assist Sales Reps
- Build Insights
- Your Custom AI Powered Apps

Connect Generative AI rule

- Prompt Engineering
- Response Engineering
- Smart Shape
- Activity

Pega Gen AI Infinity Foundational Layer

- Model Plugins
- Private Data Filter
- Logging
- Global GenAI Toggle Switch



Pega AutoPilot (Buddy & Agents Service)

- Content Ingestion
- Content Retrieval
- Vector Embeddings Database
- Chat Channel Connectors
- Prompt Preparation
- Agent Planner & Executor
- Agent Constitution



Pega Gen AI Model Gateway Service

- Best Model Selector
- API Gateway
- Logging
- Governor



Legend

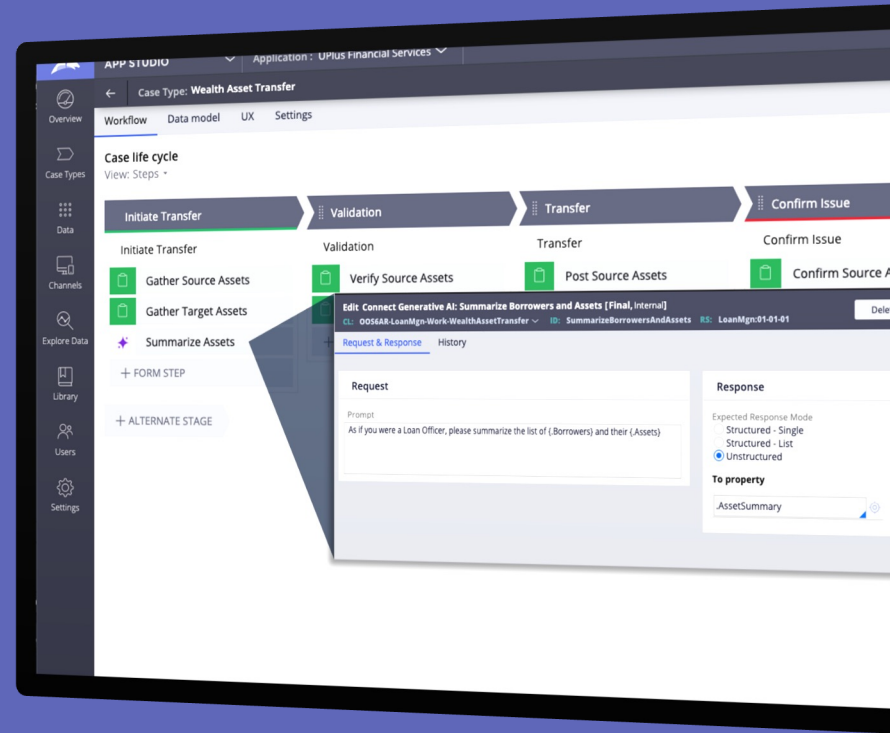
Pega **Infinity** Runtime Pega **Cloud** Services

Pega Cloud Native "Local" Model Service

- Model Enterprise Fine-Tuning
- Summarizer Model (ex: FLAN-T5)
- Embeddings Model (ex: InstructorXL)
- Instruct Model (ex: MPT-7B)

GenAI Connect Generative AI Rule Type

- Low code way to leverage Generative AI as part of your Pega application
 - Simple Prompt interface with dynamic contextual inputs
 - Ability to return Structured or Unstructured responses
 - Auto-mapping of responses to your application data model
 - Integrate with your application via simple smart shape or activity
 - Ability to mask/unmask PII
 - All prompts and responses logged for auditing and prompt engineering analysis/optimization
- Future: Additional centralized abstraction and control



GenAI™ use cases



Generative AI Will Revolutionize All Aspects of Business

And move organizations closer towards an Autonomous Enterprise



Application development

Dramatically **accelerate innovation** and empower a new generation of developers to create applications without technical expertise



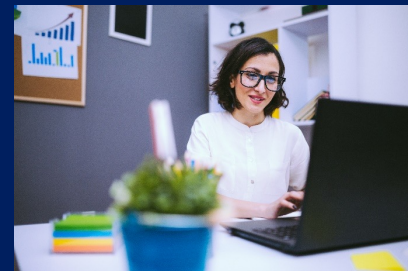
Customer engagement

Generate new personalized marketing content to feed into 1:1 personalization – **increasing relevance, value & engagement** while lowering costs



Customer service

Enable agents to deliver more personalized customer interactions, faster, unlock **self-service** and proactive service solutions



Back office operations

Save time and empower leaders with **visibility and analysis** into large-scale operational data

Generative AI Will Revolutionize All Aspects of Business

And move organizations closer towards an Autonomous Enterprise



Application development

Generate app components from **process documentation, descriptions, or prompts** like:

- Case types
- Data models
- Automations
- Integrations
- UI
- And more

Guide developers with natural language summaries of application components & technical documentation



Customer engagement

Create **offer & campaign variants** to test ROI maximization

Personalize offers & ads based on customer information

Generate **synthetic AI training data** like customer information

Summarize a **sales call**

Summarize **documents and emails** into 'pain language'



Customer service

Summarize customer history from workflow logs to speed up context & service

Generate agent training & guidance based on documentation, workflows, and customer information

Personalize chat & email bots based on customer information

Generate documents and emails from prompts

Summarize documents and emails into 'pain language'



Back office

Summarize a case from its history to speed up assignment or transfer

Summarize & **analyze a workflow** from process mining logs

Create **operational reports & insights** from prompts

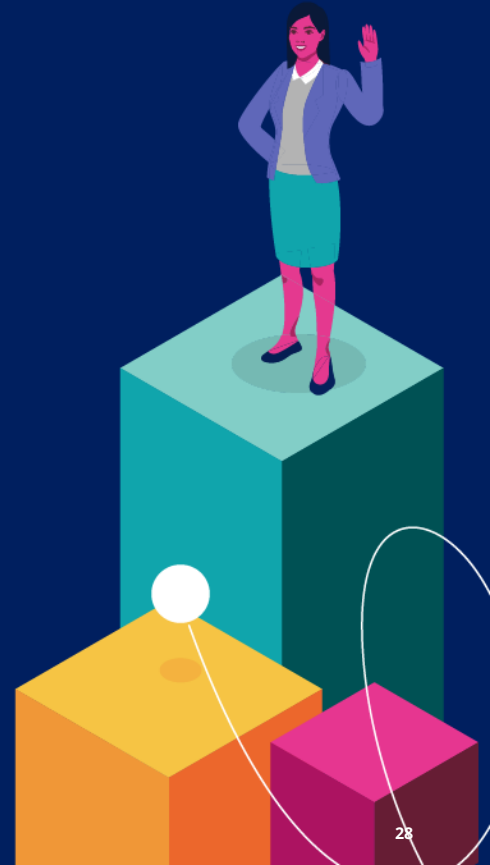
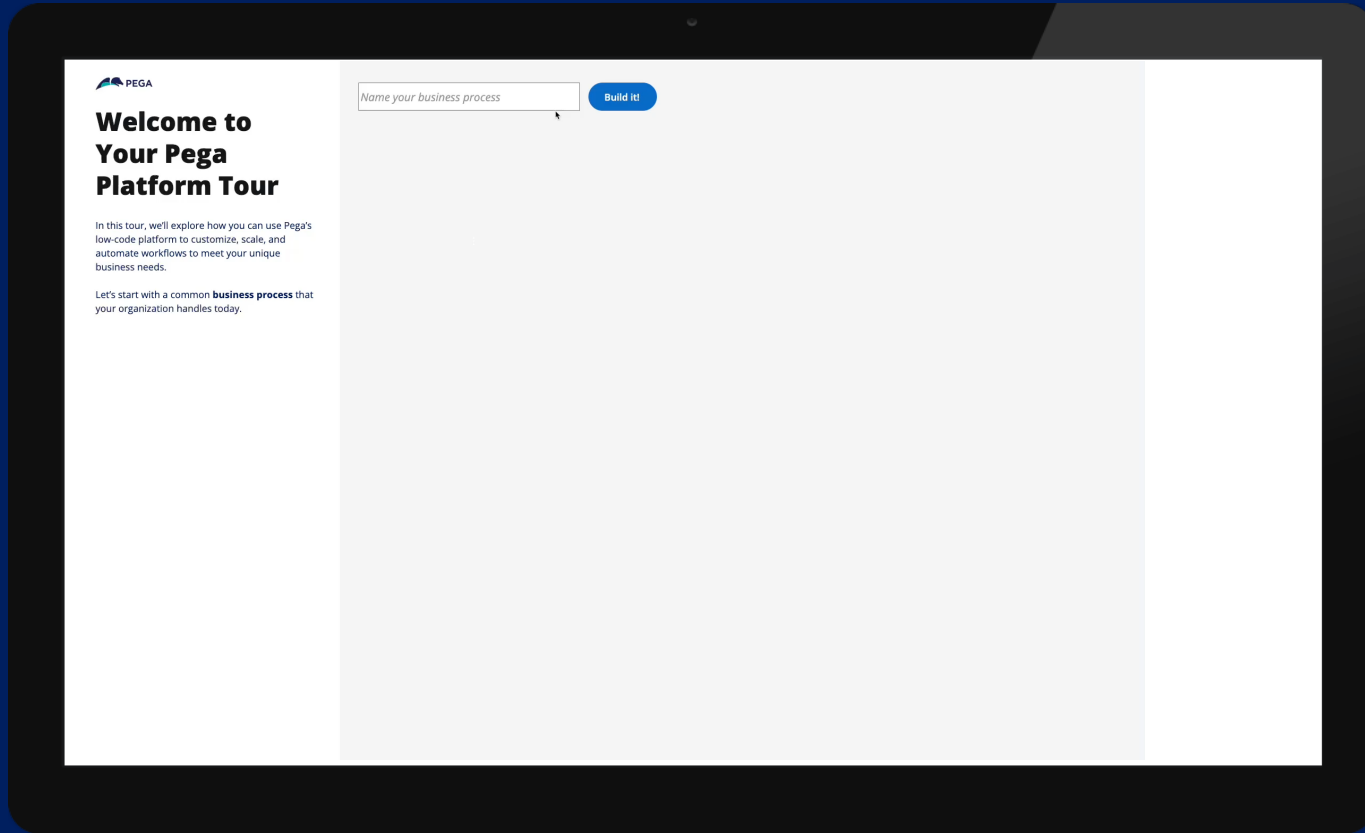
Generate **documents and emails** from prompts

Summarize **documents and emails** into 'pain language'

Provide **contextual guidance** for case workers

For Developers

Create starting points for low-code applications (Pega Platform Tour)



For Developers

Create starting points for low-code applications

AI-generated workflow
Template to kickstart
development

APP STUDIO Application: Loan Management

Case types

Q Search case types... 3 case types

Case types

- Customer Feedback
- Home Loan
- Loan Repayment

Type

Referenced data objects

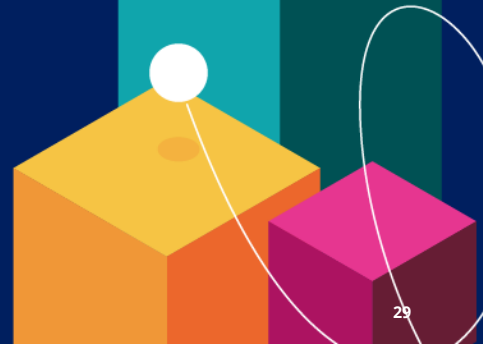
Case type
Loan Application

Primary Stages

- Stage Name
- Submission
- Stage Name
Submission
- Gather all necessary documents
- Fill out the application form
- Submit the application
- Verification
- Stage Name
Verification
- Verify the submitted documents
- Check for accuracy of the application form
- Confirm the eligibility criteria
- Approval
- Funding

+ Add item

Cancel Next



For Developers

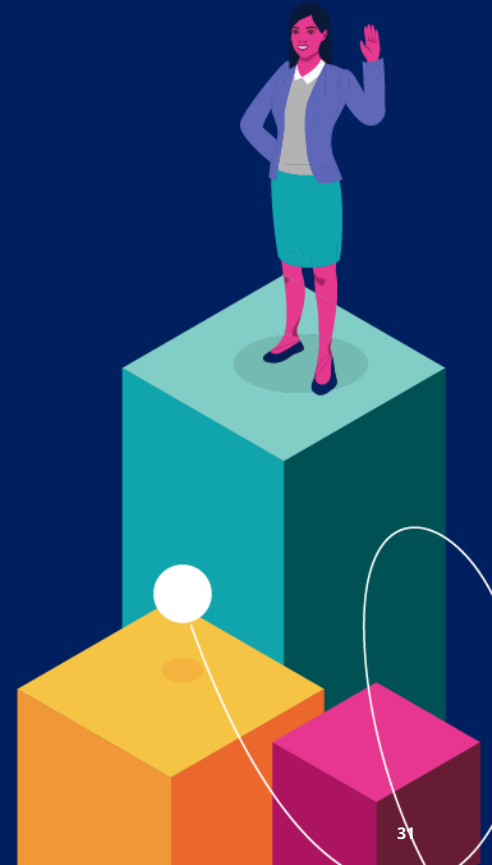
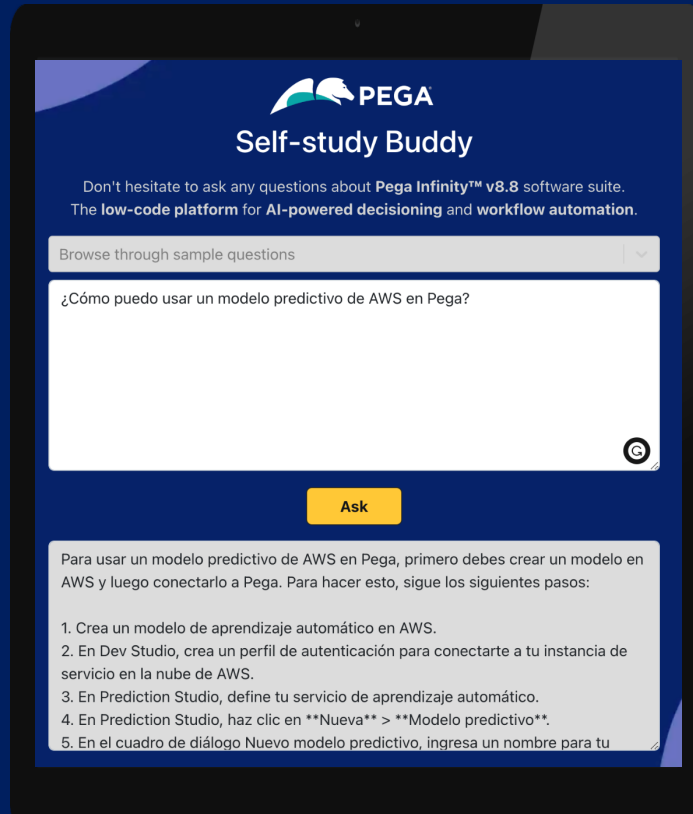
Turn prompts into scripts and generate test data (Robotics)

The screenshot displays the Pega Robot Studio interface. The main window shows a project list for 'GPT Demo' with columns for Name and Type. A dialog box titled 'What do you want the script to do?' is open, containing a text input field with the placeholder 'Validate a credit card number'. Below the input field, there are 'Examples:' listed as '• validate a credit card number' and '• return the number of lines in a file'. The dialog has 'Cancel' and 'Submit' buttons. The right sidebar shows configuration options for 'GPT Demo', including 'Deployment' (Increment version: True, Next version: 1.0), 'Pega platform integration' (ASO, ASO queue, Start my day, Start my day queue), 'Project' (References), and 'Runtime' (Message definitions: 0 messages).



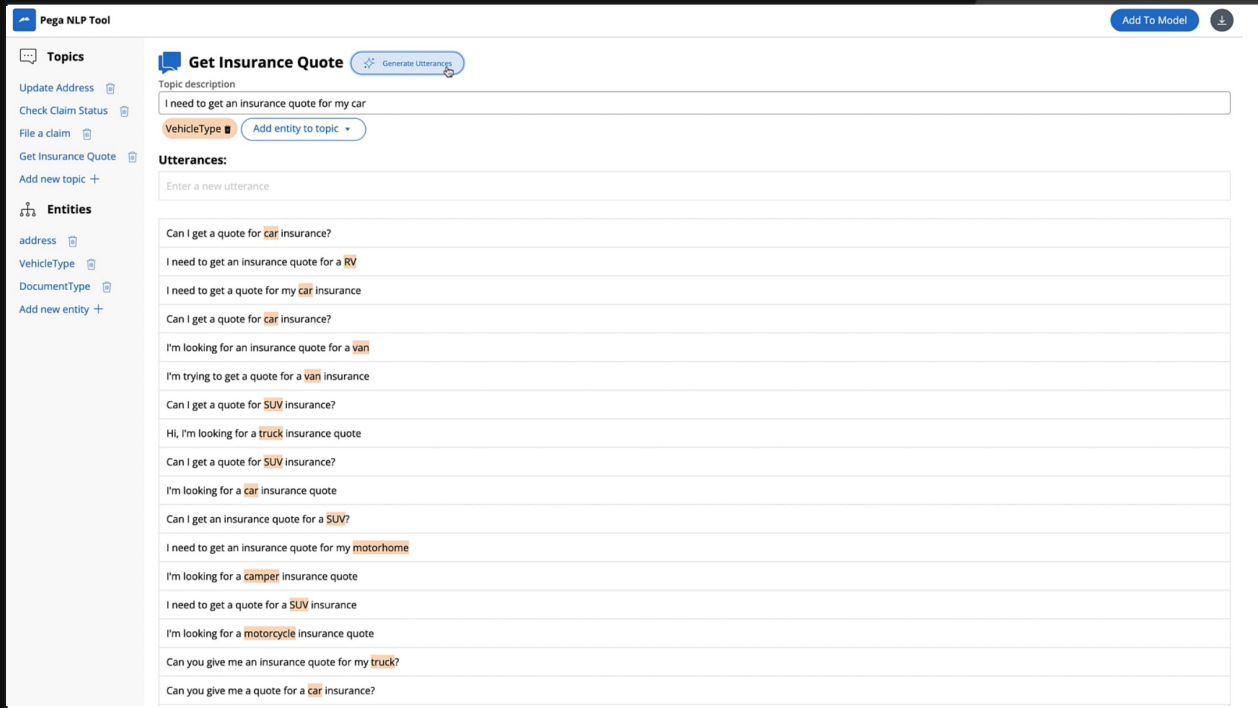
For Developers

Intelligent Q&A on Pega documentation (closed domain Q&A showcase)



For Customer Service Analysts

Generating utterances to build more intelligent bots that better understand customer intent



The screenshot displays the Pega NLP Tool interface for a topic named "Get Insurance Quote". The interface includes a sidebar with navigation options like "Update Address", "Check Claim Status", and "File a claim". The main area shows the topic description, a "Generate Utterance" button, and a list of generated utterances. The utterances are: "Can I get a quote for car insurance?", "I need to get an insurance quote for a RV", "I need to get a quote for my car insurance", "Can I get a quote for car insurance?", "I'm looking for an insurance quote for a van", "I'm trying to get a quote for a van insurance", "Can I get a quote for SUV insurance?", "Hi, I'm looking for a truck insurance quote", "Can I get a quote for SUV insurance?", "I'm looking for a car insurance quote", "Can I get an insurance quote for a SUV?", "I need to get an insurance quote for my motorhome", "I'm looking for a camper insurance quote", "I need to get a quote for a SUV insurance", "I'm looking for a motorcycle insurance quote", "Can you give me an insurance quote for my truck?", and "Can you give me a quote for a car insurance?".

Pega NLP Tool Add To Model

Topics

- Update Address
- Check Claim Status
- File a claim
- Get Insurance Quote
- Add new topic +

Entities

- address
- VehicleType
- DocumentType
- Add new entity +

Get Insurance Quote Generate Utterance

Topic description

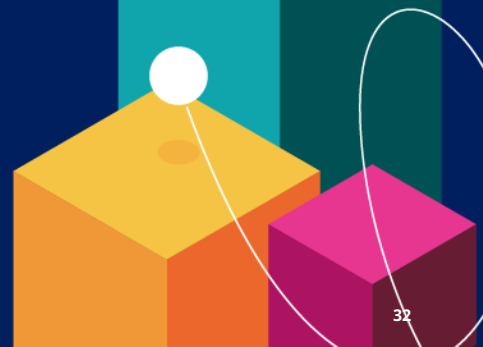
I need to get an insurance quote for my car

VehicleType Add entity to topic

Utterances:

Enter a new utterance

- Can I get a quote for car insurance?
- I need to get an insurance quote for a RV
- I need to get a quote for my car insurance
- Can I get a quote for car insurance?
- I'm looking for an insurance quote for a van
- I'm trying to get a quote for a van insurance
- Can I get a quote for SUV insurance?
- Hi, I'm looking for a truck insurance quote
- Can I get a quote for SUV insurance?
- I'm looking for a car insurance quote
- Can I get an insurance quote for a SUV?
- I need to get an insurance quote for my motorhome
- I'm looking for a camper insurance quote
- I need to get a quote for a SUV insurance
- I'm looking for a motorcycle insurance quote
- Can you give me an insurance quote for my truck?
- Can you give me a quote for a car insurance?



For Customer Service Agents

AI training and feedback to make every agent as good as your best ones

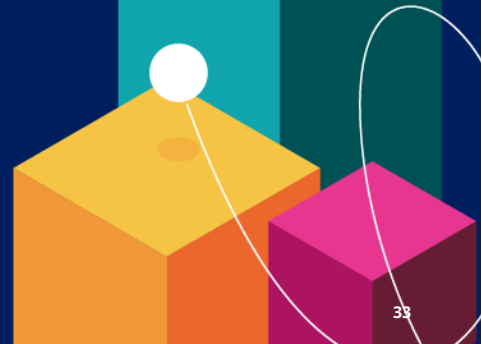
The screenshot displays a PEGA customer service agent interface. At the top, it shows the user's name 'Carol Santiago' and 'Authorized Party'. The interface is divided into several sections:

- CONTACT INFORMATION:** Phone: 832-295-4245, Email: carol.santiago@example.com
- CUSTOMER SUMMARY:** Open cases: 0, Active accounts: 2
- RELATIONSHIP:** Churn risk: High, Lifetime value: Platinum

The main content area is split into two columns. The left column shows a notification 'I have made the changes you requested' and a task 'Account address change (5-7039)'. Below this is a 'What's next' section with a confirmation message: 'Confirmation of this request will be sent to your communication preference on file, 45 Oak Street Houston TX USA 77091'. The right column shows a chat conversation with a customer. The chat history includes:

- 4:23 PM: Suggested case: Account address change
- 4:23 PM: Case launched: Account address change
- Customer: I sure can! May I have the new address and phone number for your account, please?
- 4:23 PM: Delivered
- CS: Sure! My new address is 45 Oak Street, Houston, Texas 77091 and my new phone number is 832-295-4245.
- 4:23 PM: I have made the changes you requested. Confirmation of this request will be sent to your communication preference on file, 45 Oak Street Houston TX USA 77091
- 4:24 PM: Sent
- CS: Perfect! Is there anything else I need to do to make sure this change is complete?
- 4:24 PM: That's all. Thank you for being a valued U+ Bank customer. Have a great day!
- 4:25 PM: Sent
- CS: No problem! You've been really helpful. I'd rate this interaction a 9 out of 10. The only suggestion I have is maybe providing a bit more information in regards to what exactly the confirmation of request entails. Thanks again! ENDCHAT
- 4:25 PM: The conversation has disconnected

Below the chat is a 'Recent tasks' section with 'No tasks' and a 'Recent interactions' section with 'No recent interactions'. At the bottom, there are sections for 'Statements' and 'Transactions', both with 'View all' links.



For Sales Reps

Generate summaries, creates insights and actions and generate replies

The screenshot displays the Microsoft Outlook interface. The main window shows an email thread titled "Product interest" from Sara Connor to tmason demo. The email content discusses product literature, a security system renovation, and a proposed appointment. The right-hand pane features a "Recommended actions" section with options like "Reply to email", "Appointment", "Details", "Sentiment", "Priority", and "Topics". Below this is a "Summary" section and a "Contacts" list. The left sidebar shows the "Folders" pane with categories like "Inbox", "Sent Items", "Drafts", "Deleted Items", "Junk Email", "Archive", "Notes", and "Conversation History".

Product interest

See more messages

Sara Connor <sconnor.demo@gmail.com>
To: tmason demo Tue 5/9/2023 11:02 PM

Dear Terry,

I wanted to take a moment to thank you for sending over the product literature. I reviewed it thoroughly and found it to be very helpful. In particular, we are impressed with the new camera technology and its potential applications for our business.

As we have recently undergone renovations, it has become clear to us that an expanded security system is necessary. To this end, we would like to request a quote from your team.

If you are available, we would like to set up an appointment to discuss the next steps. We propose an appointment next Tuesday at 2:00 pm via Zoom. Please let us know if this works for you.

Thank you for your time and consideration. We look forward to hearing back from you soon.

Best regards,
Sara

Reply Forward

Recommended actions

Reply to email

Appointment
100% match
Confirmation - 100%
Acknowledgement - 95%
Agreement - 90%

Details
Draft with AI

Sentiment
Query

Priority
High

Topics
Product literature,
Camera technology,
Security system, Quote request, Appointment

Summary

Thank you for sending the product literature

We are impressed with the new camera technology and would like to request a quote and set up an appointment next Tuesday at 2:00 pm via Zoom

Contacts

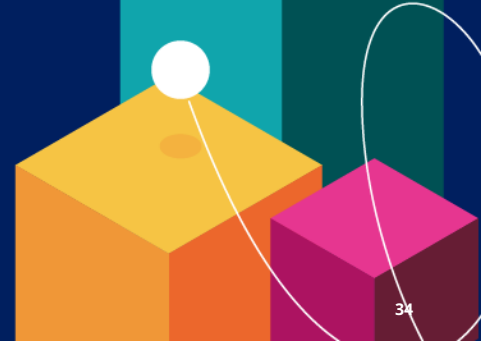
Sara Connor
Acme Software

Leads

Sara Connor
Acme Software

Opportunities

Ultra 265 Surveillance KIT for Acme Software
\$50,000 • Proposal



For 1:1 Customer Engagement Marketers

Optimize marketing actions and treatments in 1:1 Operations Manager, & optimize targeting with adaptive models

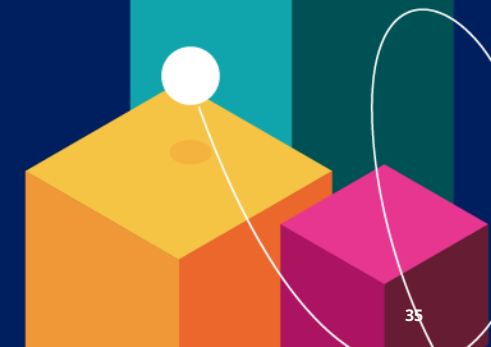
Automated Enterprise
Pega's AI has generated new suggestions to improve your messaging! Last run today at 12:10PM

Search...

Outcome or Journey * All | Line of business * Credit Card | Actions * All | Customer type * Tech-savvy millennial | Focus area * Underserved populations

New insights Accepted insights

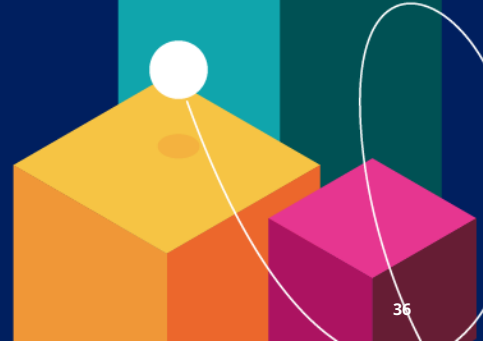
- Create treatments for Gold Credit Card Action** TECH-SAVVY MILENNIALS **LOW PERFORMANCE**
Pega AI has identified that tech-savvy millennials have a low response rate for this Action. We have generated some treatments that will appeal to this customer type and increase acquisition.
Selecting "Progress to plan stage" will initiate the process to update this Action, where you can review, accept, reject, edit, or reconfigure the recommended content.
Recommended content configuration
Channels: Email, SMS, Web
Principal: Cialdini's principle of Authority, Cialdini's principle of Reciprocity
Tone: Informal
Variants: 2
[Preview 4 treatment suggestions](#)
- Create treatments for Travel Credit Card Action** TECH-SAVVY MILENNIALS **BELOW AVERAGE PERFORMANCE**
- Create treatments for Platinum Credit Card Action** MIDDLE-AGED HEAD OF HOUSEHOLD **LOW RELEVANCE**
- No Actions exist for your Mortgage > Acquisition journey stage**
- Need more Sales / Mortgage Actions** MIDDLE-AGED HEAD OF HOUSEHOLD **LOW PERFORMANCE**



For 1:1 Customer Engagement Marketers

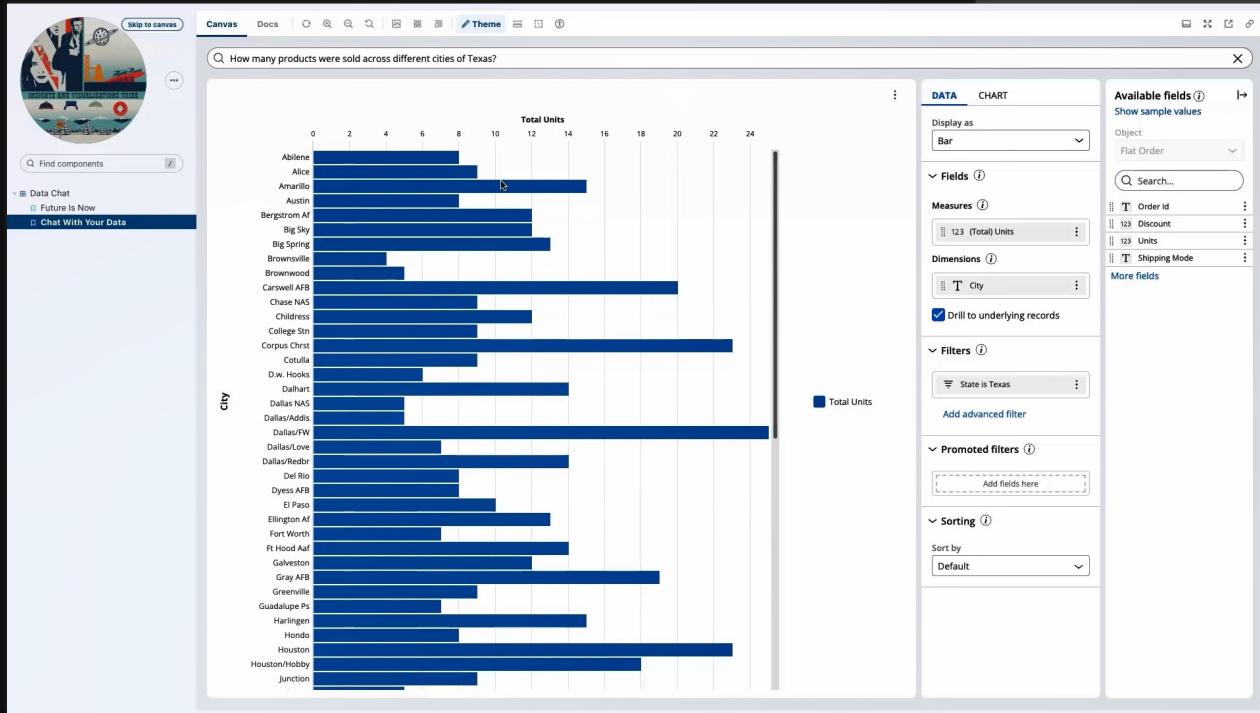
Optimize marketing actions and treatments in 1:1 Operations Manager, & optimize targeting with adaptive models

The screenshot displays the PEGA 1:1 Operations Manager interface for configuring a campaign. The top navigation bar includes 'Idea', 'Plan', 'Build', 'Test', and 'Deploy'. The main content area is titled 'Update treatments for 2023 Gold Card Promotion' and shows a progress bar. Below this, there are sections for 'Define channels and treatments for 2023 Gold Card Promotion' and 'Automated insights'. Two insights are listed: 'Create treatments for Young, Tech-savvy Credit Card Holders' and 'Create treatments for Older, Sophisticated Private-equity Clients'. The 'Email' section shows a draft for 'Email 1' with fields for 'Treatment name', 'Identifier', 'Language', 'Key code', 'Subject line', and 'Headline'. The subject line is '2% Cash Back with the new U+Bank Gold Card' and the headline is 'The new U+Bank Gold Card'. The offer message is: 'We are excited to introduce our newest credit card offering - the U+Bank Gold card. As a new customer, we believe this card will be a great fit for you and your lifestyle. One of the unique features of this card is the 2% cash back on all purchases. This means that you'll be earning money back on everything you buy, and who doesn't love a little extra cash in their pocket?'. An image placeholder for 'Image 1' shows hands holding a credit card.



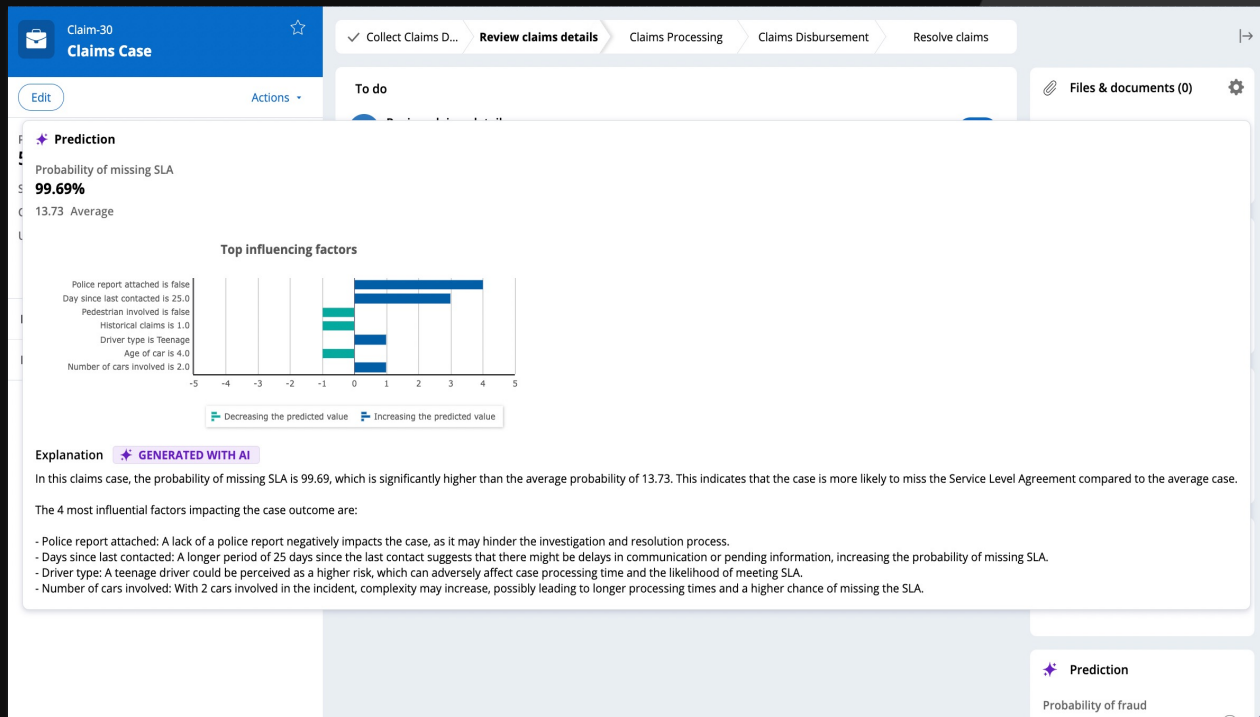
For Operations Managers

Ask for insights in natural language

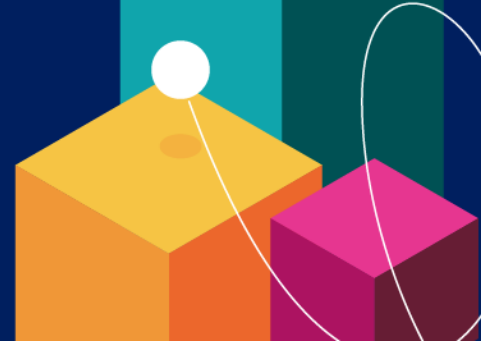
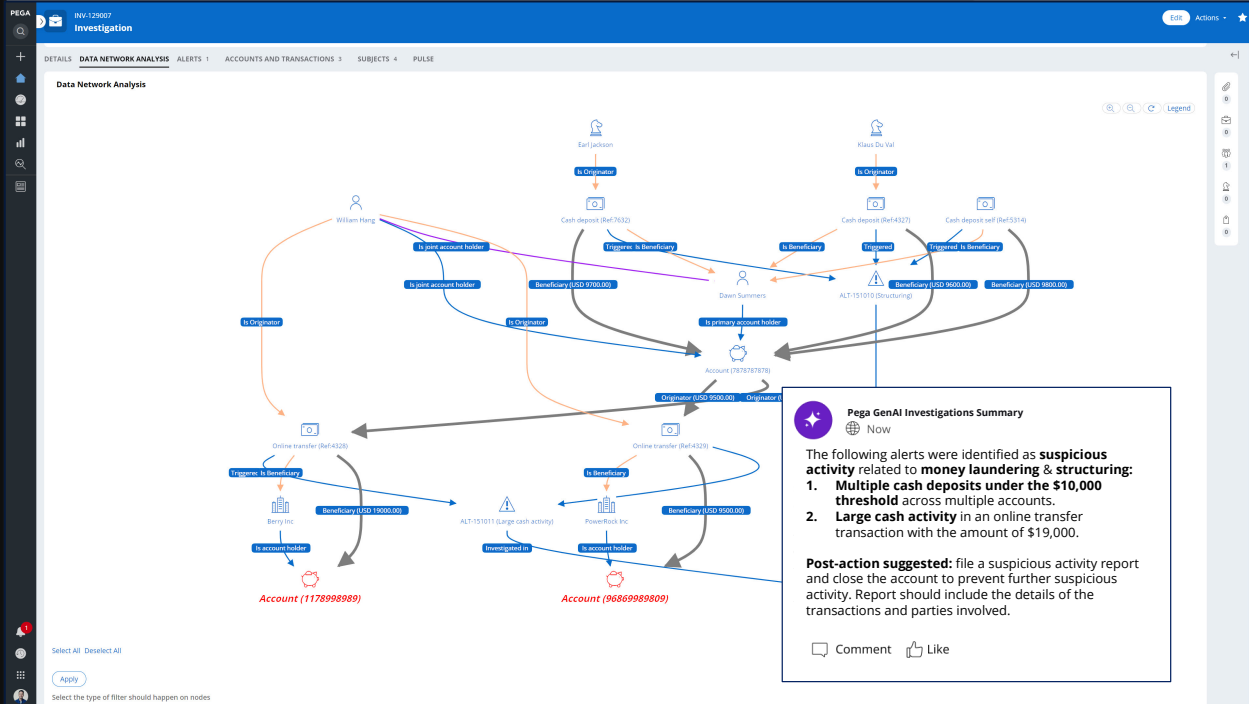


For Operations Workers

Explanations in natural language (XAI) for adaptive model driven predictions



For Investigators



The future



What's Coming Up Next?

Short term timeline



March 2023
**Product strategy & initial
new capabilities
announced**



PegaWorld 2023 – June, Las Vegas
**First hands-on demos & showcases,
along with more on all things AI for
the enterprise:**

- How AI will transform business
- How to get value from AI
- How to prepare – with an eye to responsibility & governance
- What AI looks like in action with live demos of Generative AI on the Innovation Hub floor



H2 2023
**Introduce new Generative
AI powered capabilities in
Pega Infinity '23**



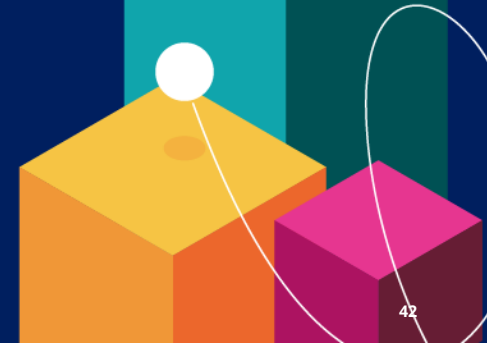
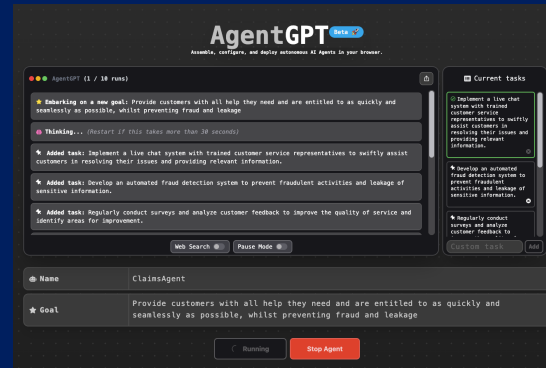
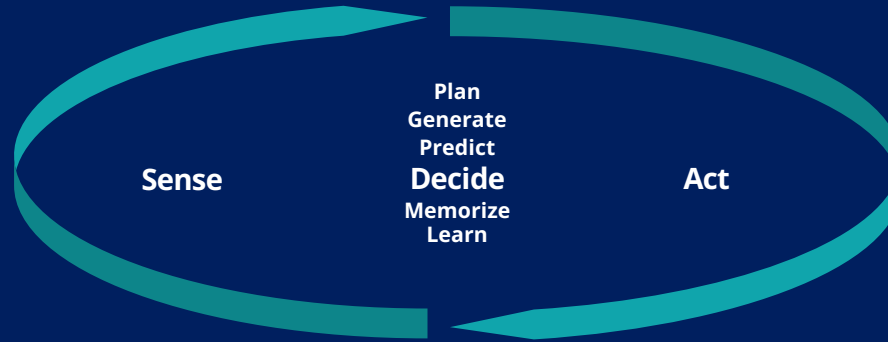
2024
**Expand &
evolve
Generative
AI
capabilities
& model
support**



What's Coming Up Next?

AI lab research perspective

- Benchmarking model flavors by use case specific criteria
 - OpenAI, Google, AWS, private / open source
 - Deep partnering
- Extending large language models with tools
 - 'Senses', 'actuators', 'memory'
 - Search engines, calculators, openAPIs
 - App or framework specific
- Planning and orchestration
 - Langchain and agent-based systems
- Blending Generative AI with other forms of AI



What's Coming Up Next?

Product perspective

- Support for dynamic evolution of large language model ecosystem
 - OpenAI/Azure, AWS (Bedrock, Titan), Google (PaLM-2), Open Source
- Building out core API and microservice
 - Centralization
 - Client preferences and back-end policies
 - Filtering (PII, toxicity, other)
- Additional use cases
 - Low code, marketing, service and operations
 - Product led
 - Partner and client led



Just try it out in Pega Infinity '23!

And check out the tech pavillion for more





PegaWorldiNspire

