

Fact Sheet – August 2018

Company Headquarters & Locations

300 Park Avenue, 12th Floor, New York, NY US
11 York Street, Level 4, Sydney, NSW 2000 Australia

Website: <https://flamingo.ai>

Firm Details

- Founded in 2014
- Listed on the Australian Securities Exchange in November 2016 (FGO)
- Primarily operates in the U.S. & Asia-Pacific markets
- To date Flamingo AI has raised US\$20M in funding

Brief Company Description

Based in NYC and Sydney, Flamingo AI is an Enterprise SaaS company in the Artificial Intelligence (AI) & Machine Learning field, providing Cognitive platforms and Conversational AI. Products include Virtual Assistants for customer sales and service and Knowledge Engines that ingest and analyse large quantities of unstructured, conversational and non-form data. Flamingo AI has significant experience in both the Financial Service and Insurance industries.

Through the deployment of Cognitive products and platforms we solve the problems of low online sales conversion rates, poor customer experience and the low efficiency of customer service.

For further information, please see this brief (two minute) [video message](#) from our CEO & founder, Dr. Catriona Wallace.

Flamingo AI's Core Product Offerings Include

1. Cognitive Virtual Sales Assistant - ROSIE
2. Cognitive Virtual Service Assistant – RILEY
3. Cognitive Virtual Enquiry Assistant - MAGGIE
4. Human Assisted Virtual Assistant - HAVA
5. Proprietary Machine Learning engine for ingesting, analyzing and optimizing large volumes of text and conversational data – LIBBY (in late stage development)

Clients

Clients include brands such as: Liberty Mutual, Nationwide Insurance, Nationwide Financial, MetLife US, MetLife Asia, Chubb, CUA, AMP, amongst others.

Highlighted Customer Use Cases

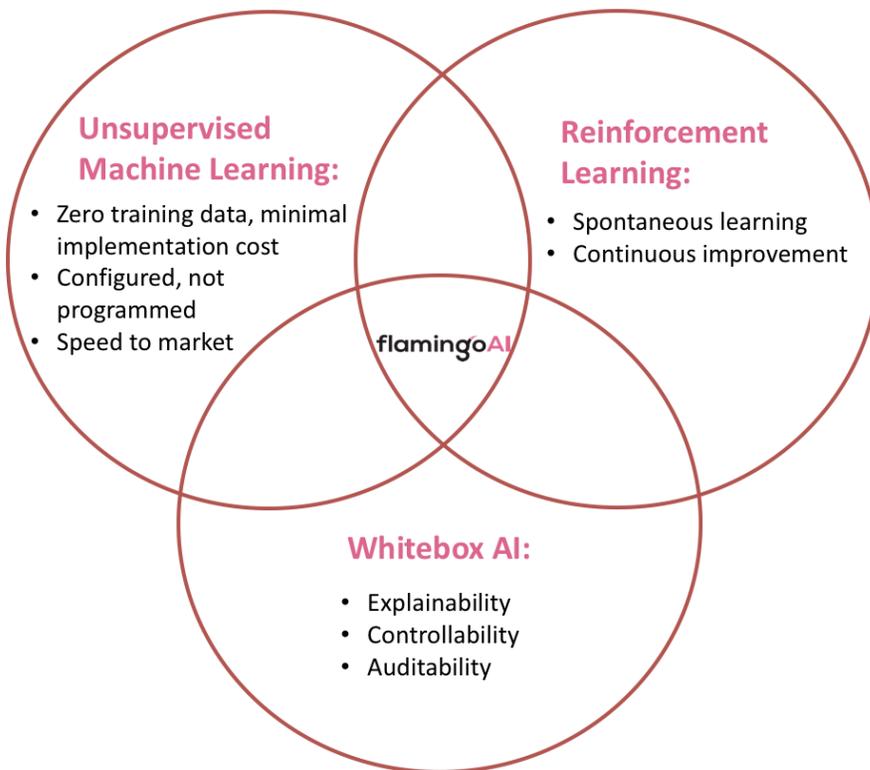
Flamingo AI's clients include multiple Fortune 100 brands in the Insurance, Finance/ Banking, and other commercial sectors. Use Cases include:

- Customer On-boarding/ Sales Journey
- Customer Renewal Journey
- Customer Service Journey

- Customer Post Insurance Claim Engagement Journey
- Employee Benefits Journey (SMB)
- Advertiser Customer Needs and Requirements – News Publisher
- Wealth Management Product Needs Journey (High Net Worth)
- Direct to Customer Retirement product sales
- Direct to Customer Retirement product – post sales service
- Auto Insurance sales
- Life Insurance sales
- Personal Loan Sales
- Travel insurance sales
- Superannuation and Life Insurance Sales
- Health Insurance Application
- General Insurance Applications
- Change of Plan
- Ingestion of email data and insights reporting
- Ingestion of conversational data and insights reporting

Core differentiators

1. Unsupervised Machine Learning with Reinforced Learning processes
2. Fast to deploy and fast to learn
3. Deep expertise in Financial Services
4. Unique IP built by Flamingo AI



IN DETAIL

ROSIE

ROSIE is a Virtual Sales & Service Assistant. “She” specializes in guiding customers through otherwise complex research, quotation, application, and payment journeys all the way through to payment.

ROSIE is also able to guide customers through the complex processes such as the filing an insurance claim, as well as through any of their service inquiries.

ROSIE can also be deployed as service assistant to internal employees and advisors/ brokers.

Potential Use Cases for ROSIE:

- Sales
- Customer Service
- Internal Service Assistant (non-customer facing)

ROSIE & HAVA

Human Assisted Virtual Assistant (HAVA) mode allows an employee to work collaboratively alongside ROSIE to guide a customer through a journey to the desired outcome. HAVA mode is particularly useful when the journey requires a personal interaction and/or personal advice.

MAGGIE

MAGGIE is built on the same technology as ROSIE, however does not need to be integrated, ie does not need to make API calls. MAGGIE is like a super smart FAQ bot or Knowledge Enquiry Engine. Use Cases for MAGGIE include:

- Concierge on client’s website
- Help facility for customers on webforms
- Internal enquiry for product knowledge – used by employees

LIBBY

LIBBY is short for Data Liberator or Self- organising Library. LIBBY ingests large quantities of unstructured, conversational or non-form data and finds patterns in these, otherwise unknown. LIBBY can become the ‘brain’ and knowledge base for a product or process. She can also determine where a Virtual Assistant is best deployed.

Technology Capabilities

Confusing web forms, and their inherent lack of interactivity, have led to low customer engagement and conversions, as well as high abandonment rates for our clients. Flamingo AI understands the complex nature of insurance and financial service products. This understanding has allowed us to create a platform that simplifies otherwise complex processes and to deliver an experience that feels intuitive and understandable to the customer. Our Virtual Assistants utilize Artificial Intelligence and Machine Learning to guide customers through various journeys, while simultaneously allowing customers to interact directly with the Assistants as they navigate these journeys via easy-to-use chat functionality.

More about Flamingo AI’s technological capabilities:

Natural Language Understanding (NLU)

Flamingo AI develops state-of-the-art Natural Language Processing technologies in-house, and deploys its very own NLU capabilities to production for use by our clients and their customers. Our core NLU capabilities include sparse/dense multi-language embeddings, sentiment analysis, topic analysis, semantic graphs, semantic space reasoning, and fuzzy NLU - just to name a few.

Most companies use third-party platforms to implement their NLU solution; they develop on top of Microsoft LUIS, Amazon Lex, or IBM Watson products. In some cases, they license the machine intelligence from another company altogether. At Flamingo AI, however, our NLU solutions are completely independent - so we can scale up easily without worrying about licensing fees or renegotiating our IP contracts every quarter. Even better; with complete control over the technology stack we can explain to our clients exactly how and why the machine learning behaves the way it does, as ROSIE's behaviour is not controlled by a third party and is therefore fully explainable.

Dialog Management

ROSIE's machine learning is Unsupervised. This means that behaviour (i.e., dialogue) is not programmed in advance. Rather, she learns very aggressively from small amounts of data. We built ROSIE to learn, rather than to be programmed or template. This allows Flamingo AI to minimize setup costs from client to client.

The other point here is subtler: we cannot foresee the future, and nor can our clients. As such, it is rarely the case that a complex business process can be known in advance. We ask our clients to predict what kinds of questions ROSIE will be asked (i.e., intents), and more often than not consumer behaviour is completely different than the intents given to us in advance by our clients. This is ultimately why dialog agents based on the Intent Identification and Slot Filling model usually fail in production. Unsupervised learning, on the other hand, lets us adapt ROSIE's behaviour in near-real-time. At Flamingo AI, we emphasize learning with our clients and discovering consumer interactions as we go forward together, rather than trying to anticipate and hard-scripting them in advance. This is the core Flamingo AI offering: a set of controlled steps into a machine learning future.

Machine Learning / NLP

ROSIE is a programming light, but learning heavy, model of Machine Learning. In conversations with Gartner, we have been told that 96% of all machine learning initiatives are essentially rule-driven systems that focus on coding up front, not learning as they go. Our central intuition is that a business process cannot be known in advance, but can be discovered by listening to what consumers say.

ROSIE implements a hybrid AI strategy embracing both unsupervised machine learning and reinforcement learning. The unsupervised nature of ROSIE allows her to start learning without any annotated training data, which massively undercuts costs and deployment lead time. ROSIE is also equipped with a reinforcement learning feedback loop, that allows her to learn continuously and spontaneously through experience. Furthermore, unlike deep learning based black-box machine learners, ROSIE's reasoning boasts full "explainability", auditability, and controllability.

Process Automation

Since ROSIE does not require any programming or knowledge engineering, clients need only periodically review what ROSIE has observed through her reinforcement learning feedback loop and approve the answers given by ROSIE or human agents. This knowledge management process is configurable on a spectrum of automation. Our client can configure ROSIE to:

- Perform aggressive learning, in which ROSIE automatically learns all observed answers without needing human approval; or
- Conservative learning, in which ROSIE remembers observed answers, but only learns the humanly approved ones; or
- A hybrid learning strategy where answers are selectively auto-approved based on certain criteria.

Although we can learn from small amounts of data, we can scale potentially out to millions of data points. The test bed for our technology includes one scenario with half a million data points that it copes with easily. We know we can go much higher than this. Because we know the technology and we have complete control over a very scalable set of algorithms, we can easily develop data mining software to automatically discover patterns in any large data set. The key to Flamingo's AI data strategy is the idea of scaling right out - we start small and get learning straight away. But no matter how large the demands of our clients grow or how many questions they answer, we know we can grow with them.

Human Involvement

To best cater to our clients' varying AI strategies, ROSIE can be deployed over a spectrum of automation: a progression from unautomated call center handling with occasional intervention, to completely automated customer service. In Auto-Pilot mode, ROSIE is capable of automatically guiding a consumer through an arbitrarily complex customer journey (e.g., sales, claim) built with our Journey Assist platform beforehand.

As soon as a customer starts to chat, Auto-Pilot will disengage and ROSIE will hand over control to an online human agent to answer customer inquiries. Human agents can later re-engage Auto-Pilot. In HAVA mode, ROSIE co-pilots with an online human agent to collaboratively guide a customer through a journey. Upon receiving a customer chat, ROSIE will perform lightning fast fuzzy semantic reasoning to identify customer intent(s) and retrieve the best, previously-seen answers (that have been approved by the client), and present these answers to the human agent, along with her confidence scores in how well those answers would fulfil the customer's information needs.

ROSIE perceives how the human agent handles the customer chat, and reinforces her belief regarding the correct answers. She then incorporates such information into her next brain build (on-demand or scheduled, depending on which mode the client had signed up for). In Full-Auto mode, ROSIE guides consumers through a customer journey and handles customer chats completely independently, without any human agent involved. Her behaviour can be fully configured through a Graphical User Interface with numerical thresholds defining her states and escalation strategies.

Analytics & Reporting

This allows us to provide customizable visual analytics that provide our clients not only the detailed statistics about how their clients interact with the platform, but also rich insights about their business processes and potentials for optimization.

The other innovative aspect of our reporting lies in communicating the maturity of the brain delivered to the client. This maturity is gained over time and is only measurable due to incremental training capability and a white box approach that we followed in building our Machine Learning technology stack.

Culture of Security

Flamingo prides itself on our Culture of Security.

Given the importance of safeguarding our clients', as well as our own, information and data, this culture pervades every level of our organization – from our executives down. It is embodied in our highly capable and trained personnel, as well as in our securely designed and built technological infrastructure, and it has allowed us to obtain the following security and compliance certifications/accreditations:

SOC2 Certification

Flamingo has achieved SOC2 Type 1 Certification, with a Type 1 Report prepared by an independent third party. The attainment of this certification demonstrates that Flamingo has comprehensive policies and procedures, as well as a securely designed and built technological infrastructure, to support the following Trust Service Principles:

- *Security*: Flamingo's systems are protected against unauthorized access (both physical and logical).
- *Availability*: Flamingo's systems are available for operation and use as committed or agreed.
- *Processing Integrity*: System processing is complete, accurate, timely, and authorized.
- *Confidentiality*: Information designed as confidential is protected as committed or agreed.

If you are interested in reviewing the SOC2 Type 1 Report, please email hello@flamingo.io. Flamingo anticipates the attainment of SOC2 Type 2 Certification in late 2018 or early 2019.

PCI Compliance

Flamingo has also attained full compliance with the Payment Card Industry Data Security Standards (PCI DSS), as a Level 4 Merchant. Achieving PCI DSS compliance certifies both that the company has sufficient processes and procedures, as well as a secure underlying technological infrastructure, to support the processing of credit card transactions.

Customer Engagement Strategy (POC – Post POC Production - Enterprise):

Flamingo AI engages clients through the following process:

Key Client Roles:

- CEO
- CMO
- CIO
- CTO
- CDO
- Head of Sales
- Head of Contact Center

Production Pilot: A Production Pilot is designed to prove out the technology in an initial use case. A fixed fee is charged for the set-up of the Production Pilot. Pilots generally take 8 weeks to set up - including design, integration, and machine training. The Pilot itself then runs live customer sessions for two months.

Scale Roll Out: Based on the success of the Pilot, the engagement can then move immediately into full production with unlimited scale.

Enterprise: Wide scale rollout across multiple use cases and lines of business.

Revenue Model:

Proof of Concept (POC): Fixed fee Production Pilot including two months of live production

Production (Post POC): Cognitive Virtual Assistant Monthly Subscription Fee (billed annually) including predefined KPI attainment

Enterprise: Multiple Virtual Assistants both sales and service (discounts available)

Key Differentiators:

- True Machine Learning
- A complete Application NOT an AI generic platform
- Designed for implementation by business users (no programming required)
- Sophisticated security practices and procedures that have enabled Flamingo AI to meet rigorous F100 company security assessments (PCI & SOC2 Type1 certifications pending)
- Customer experience and sales conversion experts
- Vertical expertise (Financial Services, Insurance, and Telecommunications)
- Machine Learning as a product that delivers commercial benefit
- Technology that is contextually sensitive and has the ability to understand and process semantics
- Completely unsupervised machine learning, requiring no data science expertise
- In-house IP: no fees, no lexica maintenance, and no syntactic pattern matching needed to scale

- Auditable and controllable learning interface ensures learned answers are compliant before being used
- Visualization tools allow clients to see inside and manipulate their AI brains
- Infrastructure-in-a-box means Flamingo can be deployed anywhere on any cloud provider
- Uses client's own single source of truth and business logic
- Modular, redundant, platform built to scale across the globe