Towards Reinforcement Learning In Revenue Management

AGIFORS: Revenue Management SG Meeting 2019
Panama, Panama City

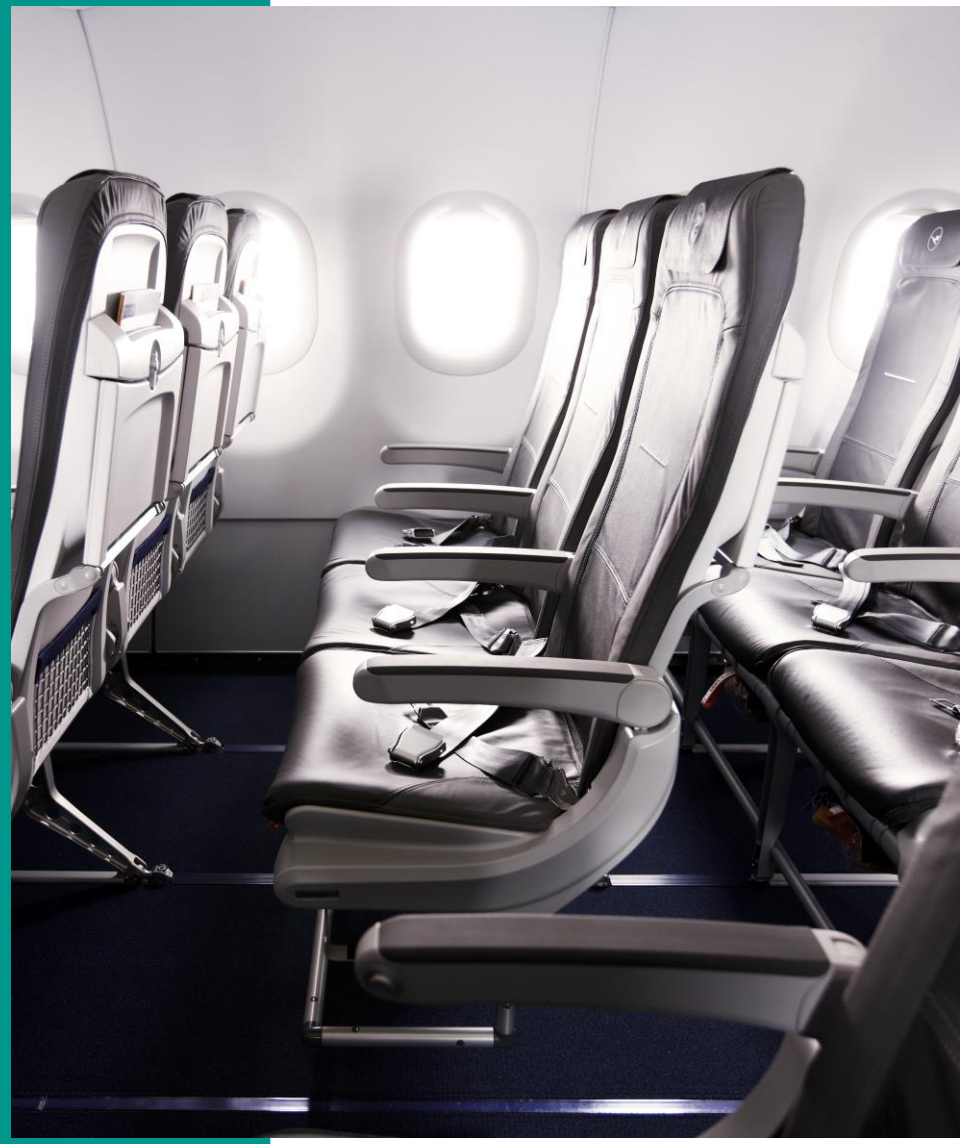

ZERO G

Combining the best out of two worlds: zeroG & Lufthansa Systems

2015

Founded by Lufthansa Systems
Dynamic start-up to keep up with technical developments

>50

People work at zeroG
With a very diverse and international background

>20

Assignments within the aviation world
Working in both commercial and operation areas

ZERO G

zeroG: a new breed of digital tech start up. We help airlines turn their data into value through advanced analytics and digitization. Our team comprises of young, dynamic & highly international technology professionals who bring along airline know-how & analytics expertise like two sides of a coin. In a nutshell, we support airlines through their digitization journey by unlocking the intrinsic power of data.

We take you on an AI journey

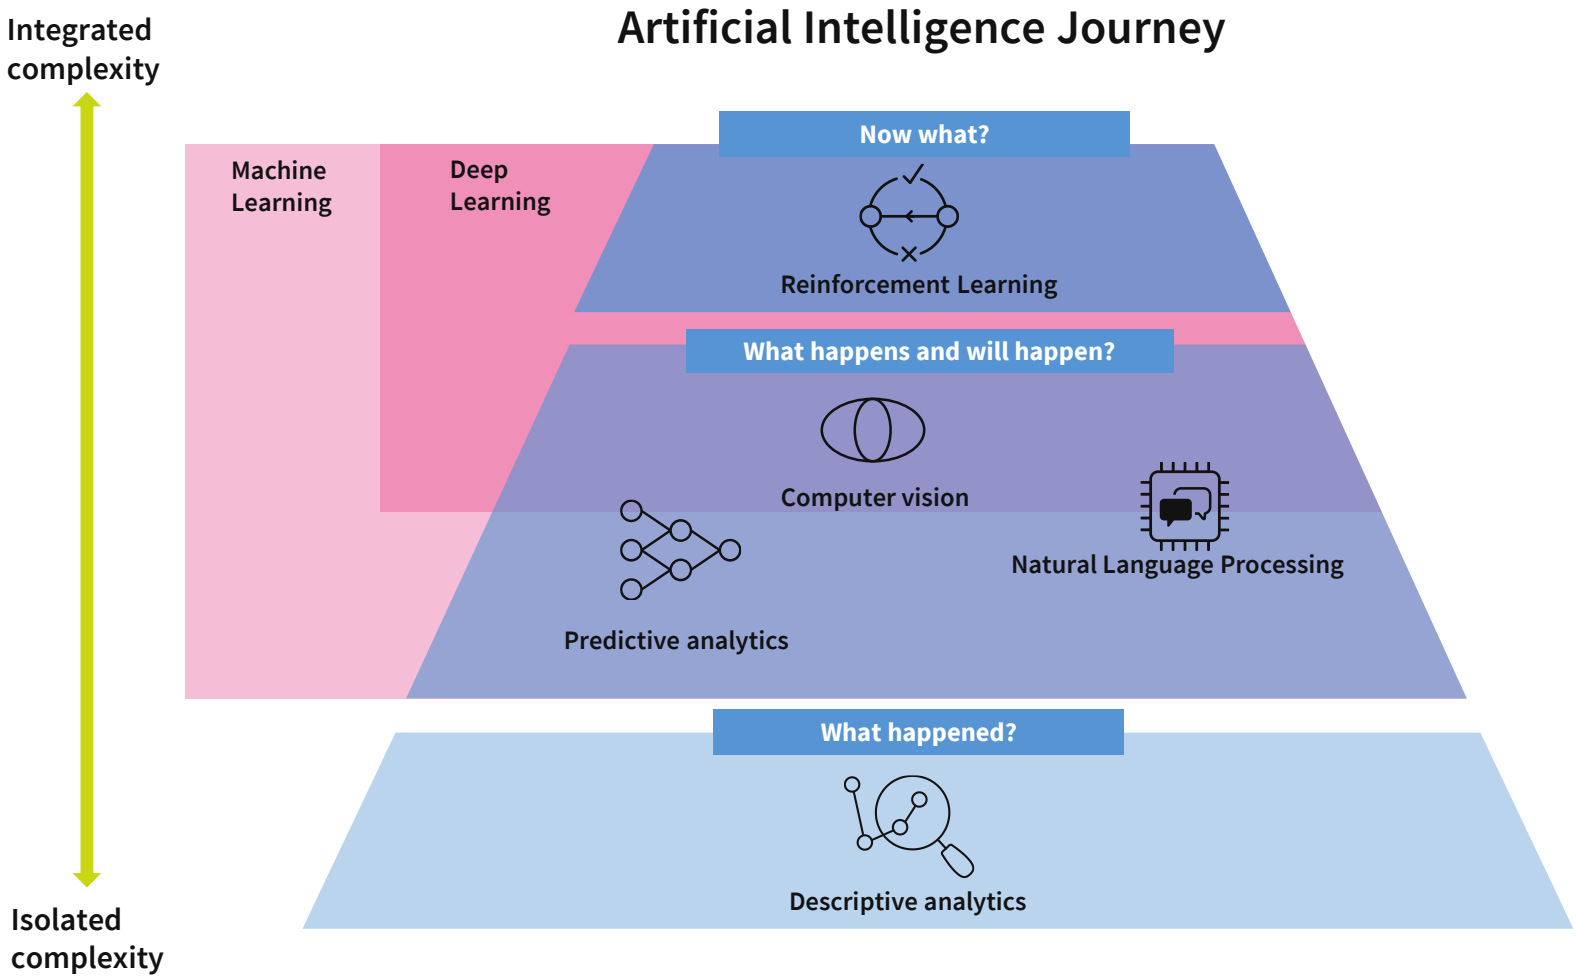

ZERO G

How does Reinforcement Learning work?

Key Elements

AI Agent
Autonomous system that will teach itself through trial and error.

Environment
World in which the agent acts, learns and aims to influence.

Observation Input
Representation of environment that serves as the basis for the agent's decisions.

KPI Driven Reward System
Essential KPI's/goals that teaches Agent how well he is performing through reward and punishment.

Situational Sketch

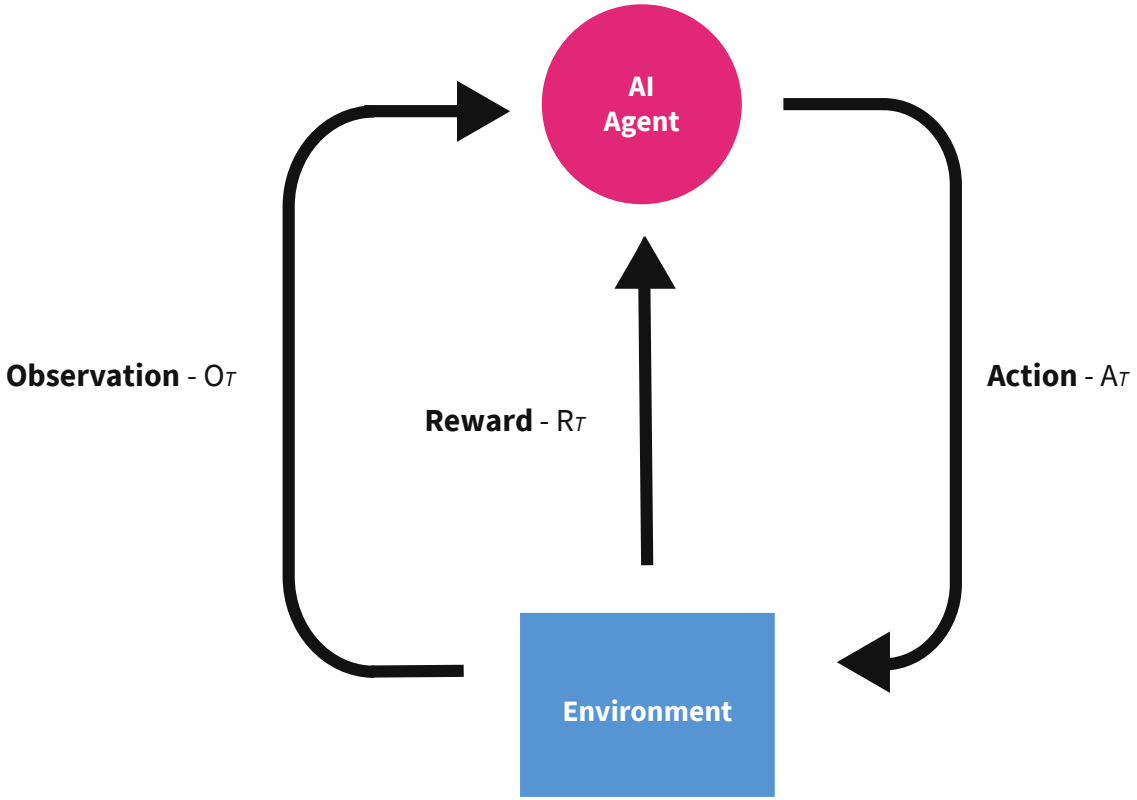

IMPACT: agent independently develops strategy to optimize actions towards highest possible result

ZERO G

Inspiration: Google Deepmind - AlphaGo

Key Elements

AI Agent
Alpha Go

Environment
Chinese game Go – the most complex game in the world!

Observation Input
Game environment, position of stones on the board, previous moves

KPI Driven Reward System
Beat opponent!

Situational Sketch

AlphaGo now beats the best human Go-Player

ZERO G

Why reinforcement learning?

01. End-to-end optimization

An AI Agent learns how to reach a long-term goal all the way from the initial actions leading to it

02. Innovation

New strategies are explored to improve continuously.

03. Data-driven Decision Making

If the environment changes, the agent will adapt its behavior accordingly.

04. Learning and Adapting

Influence of different features of a situation is learned directly from the data.

Project Deepsky

A playful approach towards Reinforcement-Learning

ZERO G

Airbucks: An Airline Business Simulator

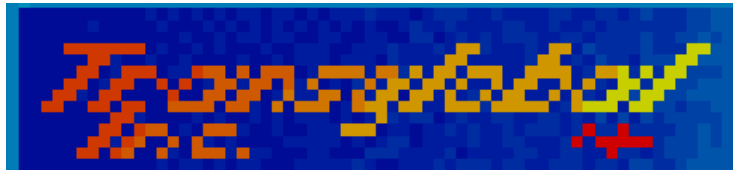

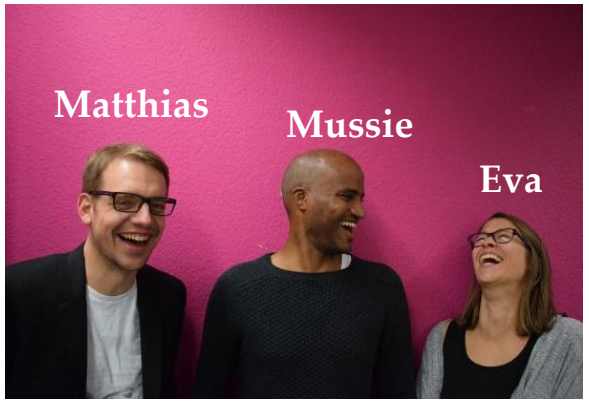

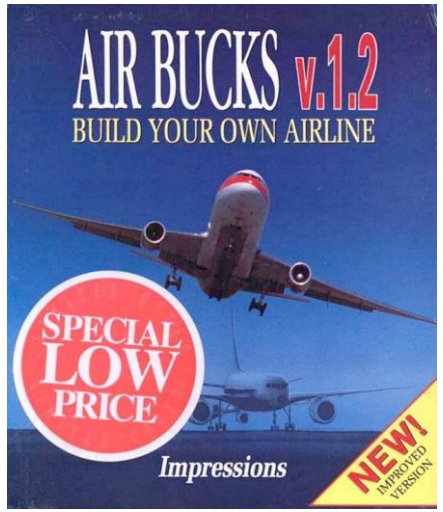

ZERO G

Result: zeroG's DeepSky

Key Elements

AI Agent

DeepSky

Environment

Air Bucks – an airline management computer game.

Observation Input

Bank balance, fleet size, landing rights, date, available airport, company worth, etc.

KPI Driven Reward System

Build an airline that will generate biggest company worth within four (game) years.

Situational Sketch

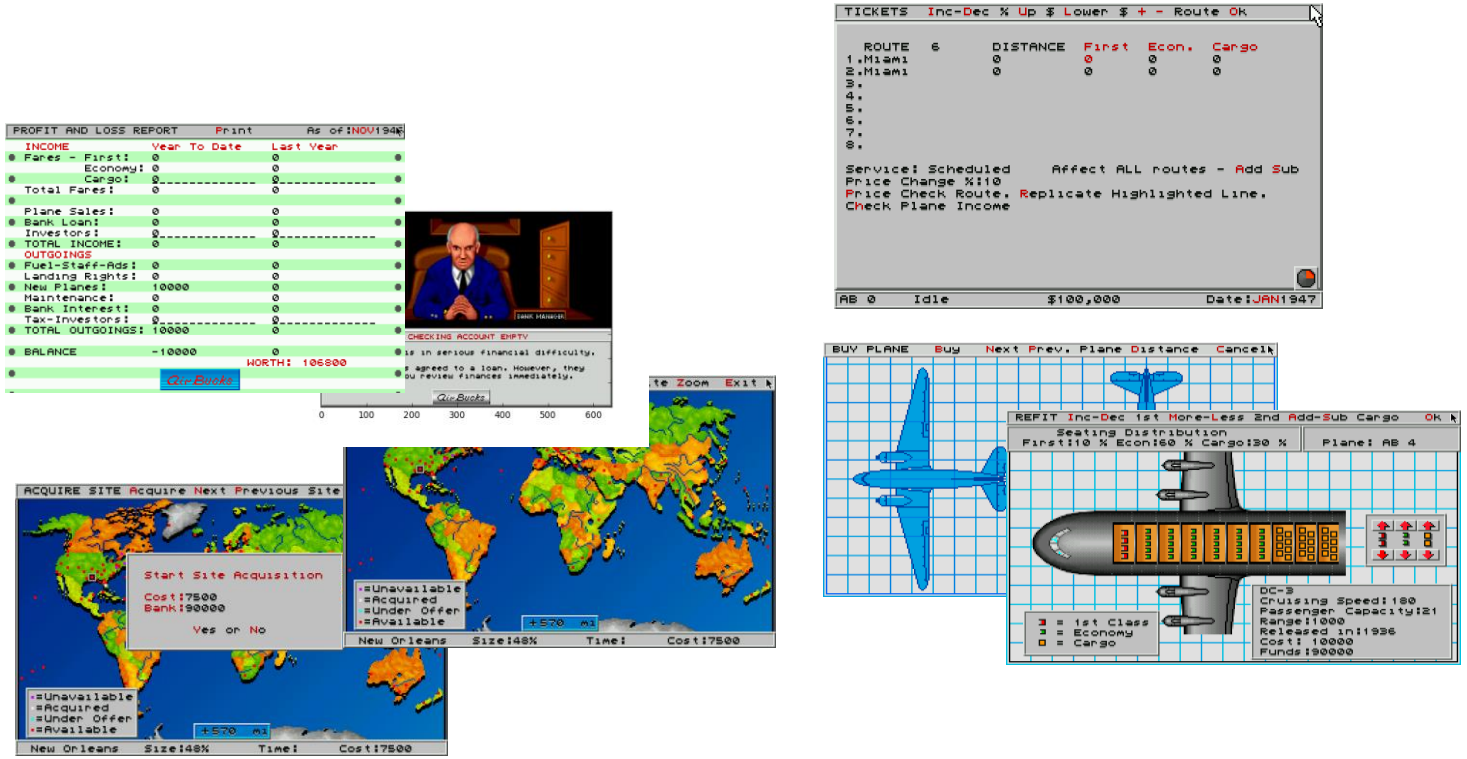

ZERO G

Result: zeroG's DeepSky

Key Elements

AI Agent
DeepSky

Environment
Air Bucks – an airline management computer game.

Observation Input
Bank balance, fleet size, landing rights, date, available airport, company worth, etc.

KPI Driven Reward System
Build an airline that will generate biggest company worth within four (game) years.

Situational Sketch

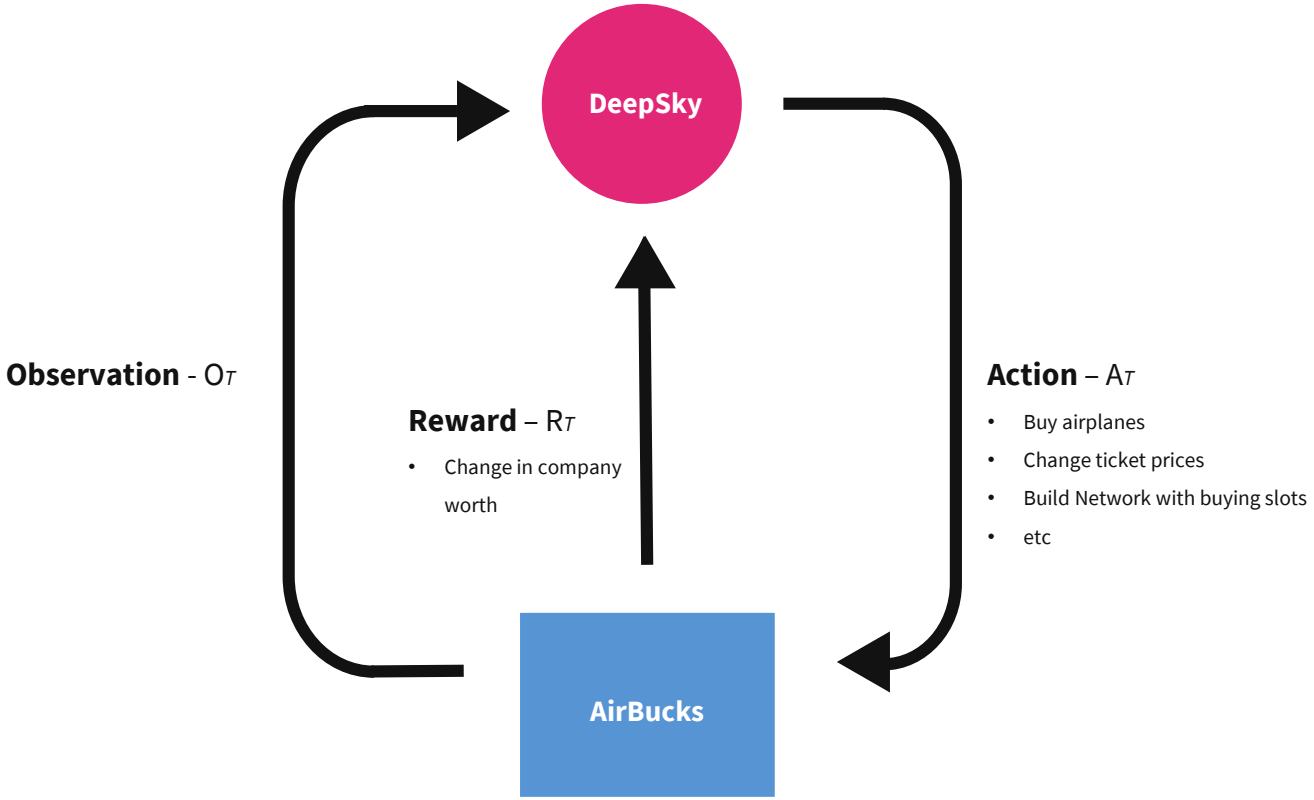

DeepSky now beats every human airline manager in Air Bucks

Our Roadmap of RL in Revenue Management: Starting from an airline simulation game, ancillaries and bundles before managing availability

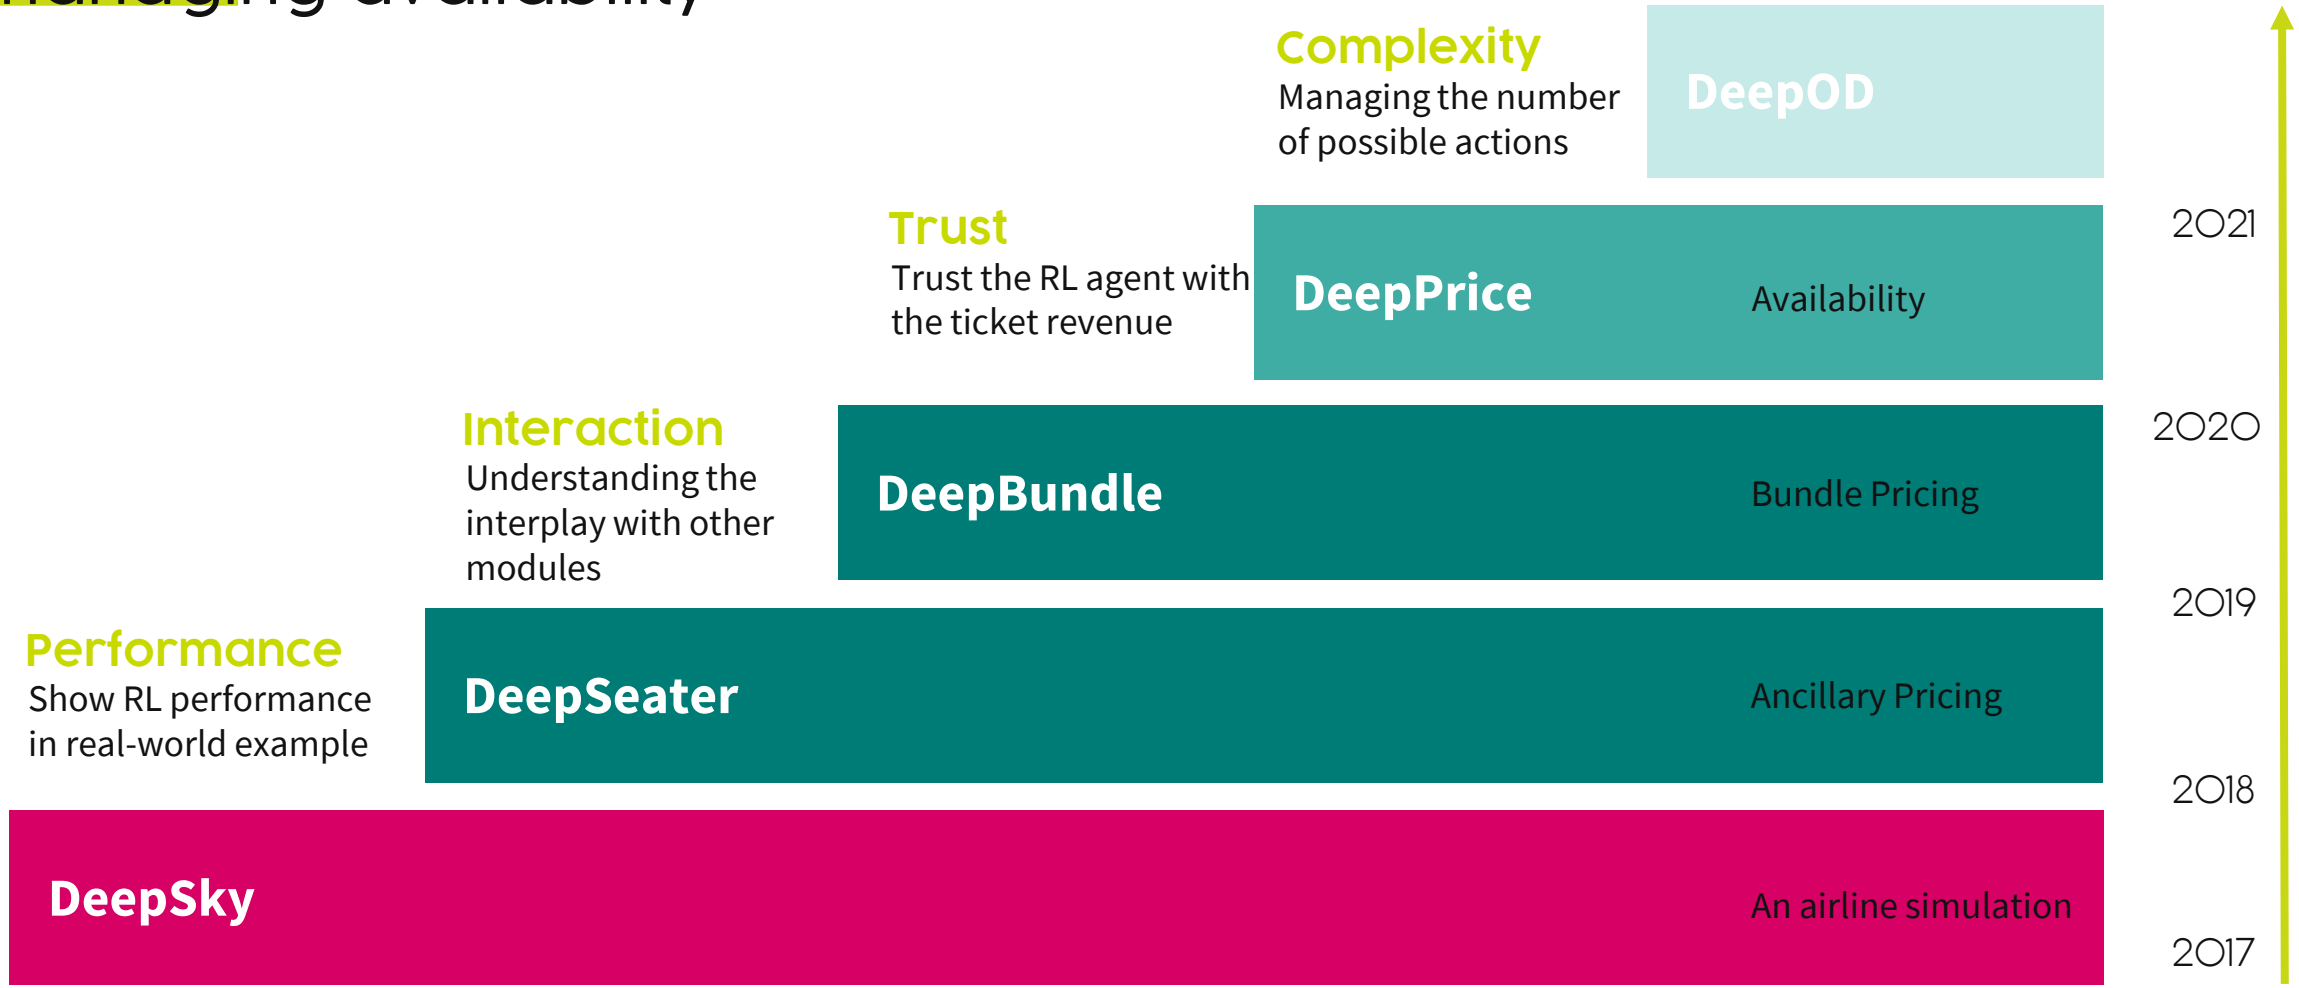

ZERO G

Dynamic Ancillary Pricing

DeepSeater

ZERO G

DeepSeater Case: AGIFORS Air, an incredibly successful LCC with standard and eco+ seats…

Situation AGIFORS Air

Static pricing for standard and eco+ seats priced at 15 USD and 30 USD respectively for all routings and departure days.

No optimization that adjusts pricing based on different demand by market, time of departure or passenger mix.

- 50 A320 aircraft with 189 seats (24 Eco+, 165 standard)
- 150 flights per day, approx. 55,000 p.a.
- Average paid Eco+ load factor of 30%
- Average paid standard load factor of 10%

➢ **Approximately 25 million USD annual seat revenue**

DeepSeater Potential

Dynamic (daily/by booking) adjustment of ancillary price on flight level (e.g. 10JAN LAX-JFK is priced independent of 11JAN LAX-JFK)

AI continuously learns and optimizes pricing for Eco+ and standard seat, adapting to changing environments.

➢ **Potential of a 10-30% increase, equivalent to 2,500,000-7,500,000 USD.**

ZERO G

Reinforcement Learning for DeepSeater requires adequate modeling of situation, action, reward

Actions

Set price point for flight AGIFORS AIR 1234 on 10MAY.

- Potential: different price point by OW/RT, length of stay, request type, bundle choice.

Situation description

Number of bookings, revenue, available seats, days before departure, passenger data

KPI Driven Reward System

Flown Eco+/standard seat revenue at departure

- Potential: Adjust reward to improve „paid“ eco+ load factor.

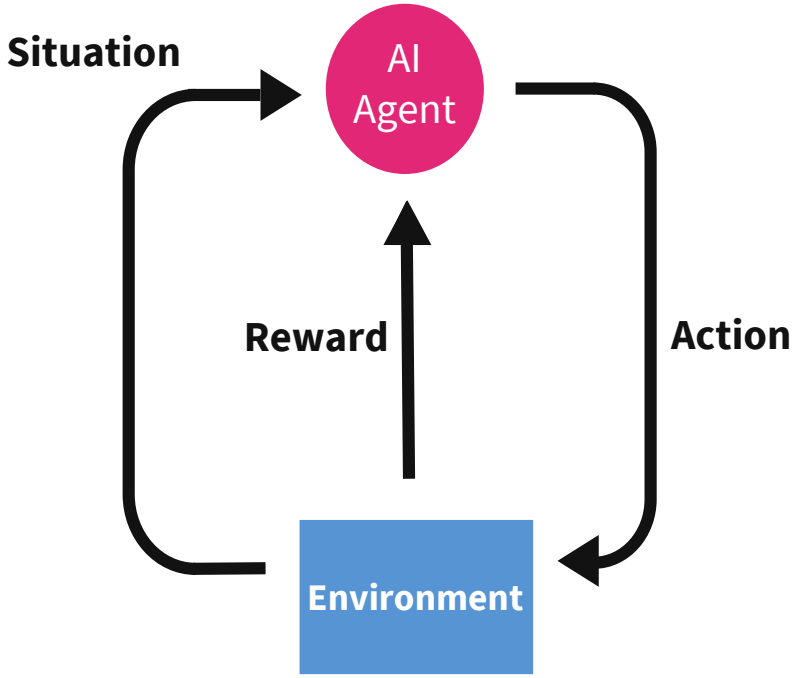

ZERO G

Dynamic Pricing for Bundles

DeepBundle

ZERO G

DeepBundle Case: More challenging but also more rewarding

Situation of Bundle/Upsell Pricing

No optimization that adjusts upsell based on different demand by market, time of departure or passenger mix.

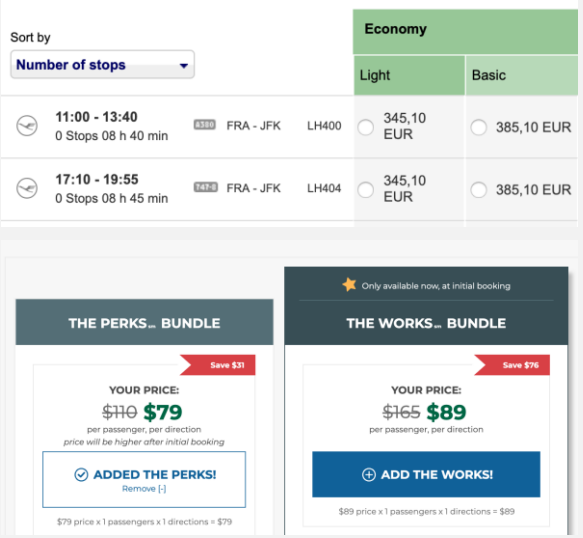

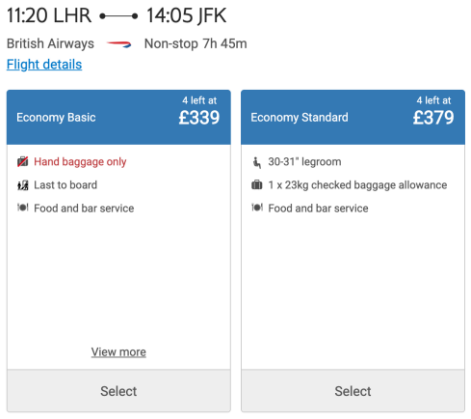

DeepBundle Potential

Adjust fare bundle price more dynamically

AI continuously learns and optimizes pricing for upsell price, adapting to changing environments, day of week trends and seasonality.

Challenges

- Competitive issue (“matching of upsell”)
- Reward definition

ZERO G

Reinforcement Learning for DeepBundle increases the complexity for the reward definition

Actions

Set price points for upsell from basic/light to economy for flight AGIFORS AIR 1234 on 10MAY.

- Potential: different price point by OW/RT, length of stay, request type, bundle choice.

Situation description

Number of bundle purchases, revenue, days before departure, passenger data, available ancillaries (if constrained)

KPI Driven Reward System

Flown upsell revenue at departure

(deducting any marginal costs of included ancillaries)

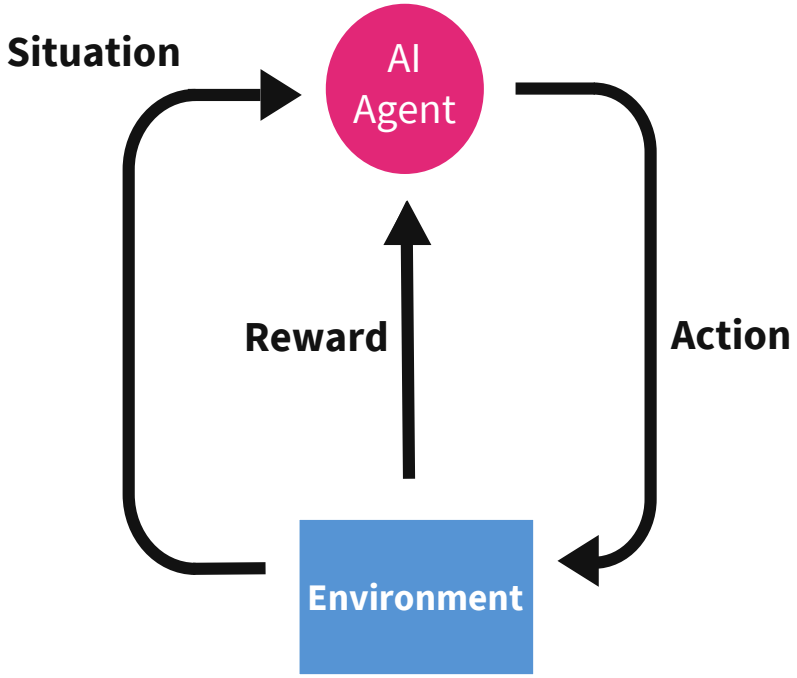

ZERO G

Dynamic Pricing for Tickets

DeepPrice

ZERO G

Reinforcement Learning for Tickets, effectively giving availability to pre-defined discrete price points on a local point-to-point flight.

Actions

Give/restrict availability for flight AGIFORS AIR 1234 on 10MAY.

- Potential: different price point by OW/RT, length of stay, request type, bundle choice.

Situation description

Number of bookings, revenue, available seats, days before departure, passenger data

KPI Driven Reward System

Flown ticket revenue at departure

(deducting any marginal costs of flying one seat)

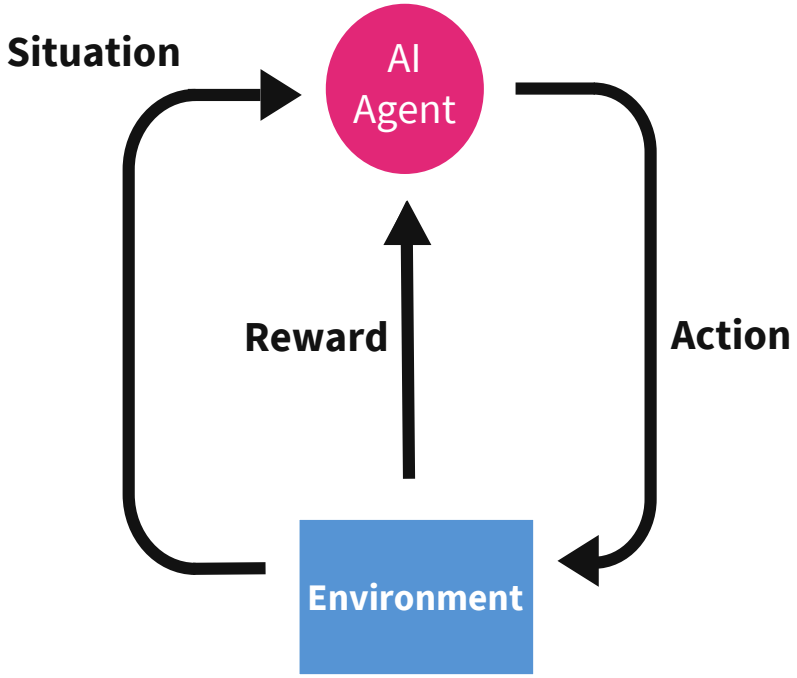

ZERO G

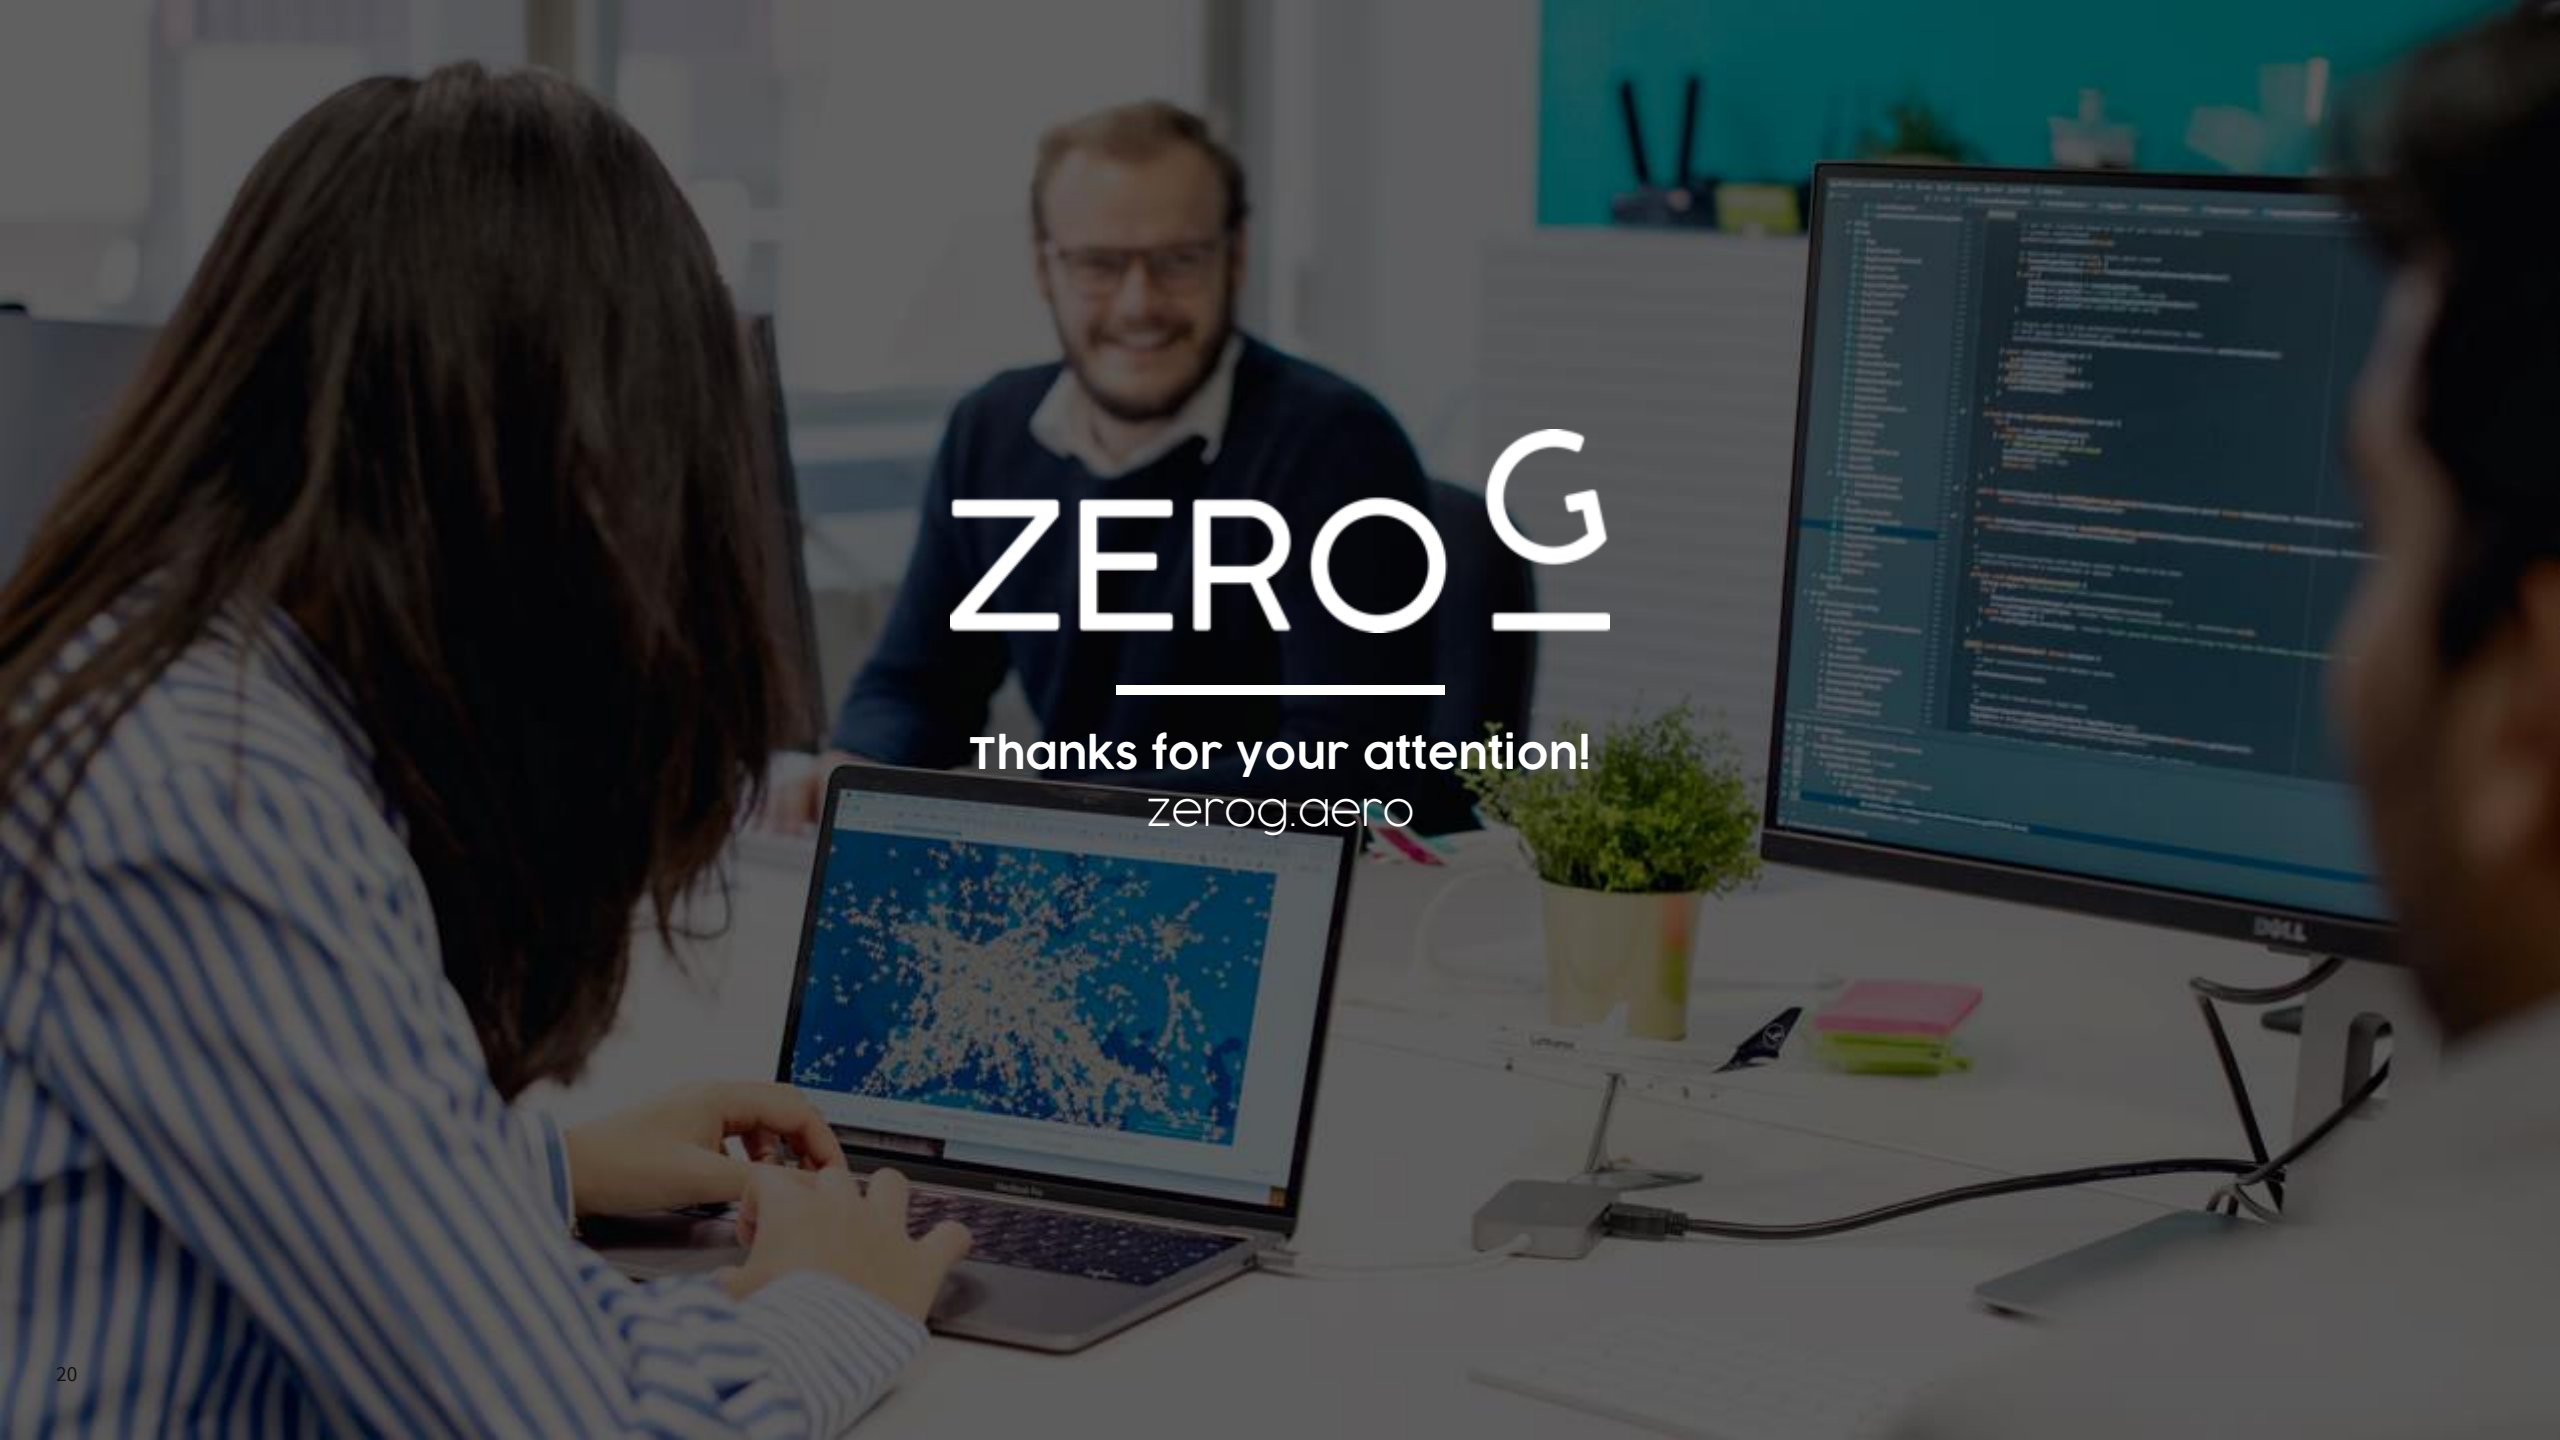

ZERO G

Thanks for your attention!

zerog.aero

Contacts

Senior Data Scientist

Mussie Beian

Mussie.Beian@zerog.aero
LinkedIn: Dr. Dr. Mussie Beian

Senior Business Data Analyst

Christian Spann

Christian.Spann@zerog.aero
Linkedin: Christian Span

ZERO G