

Data Scientist, ASX Data Analytics

ASX Position Description

ASX “All Roles Flexible”		
<p>ASX offers mutually beneficial flexible working arrangements. We recognise that employees need to balance work and personal lives.</p>		
ASX Opportunity Snapshot		What’s On Offer
Role:	Data Scientist	<p>How would you like the opportunity to work with some of the best financial markets data in the country while creating commercial data solutions for high-profile customers?</p> <p>ASX has established an ambitious new business line to commercialise its rich datasets through a data analytics platform that will sit at the centre of a financial markets data ecosystem.</p> <p>An exciting opportunity exists for an experienced Data Scientist to join this dynamic startup within one of Australia’s most recognisable and respected brands.</p> <p>This key role will lead the data science capabilities that contribute to the development of analytics-based data products as well as building the data science and machine learning capabilities of the platform.</p>
Reports to:	Business Lead, ASX Data Analytics	
People Management:	N/A	
Budget:	N/A	
Team:	ASX Data Analytics	
Date:	October 2018	
Location:	Sydney, NSW	
Flexible Role:	Yes	

What you’ll do:

- Mature ASX to be a leader in applying advanced analytics, machine learning and AI models to solve real problems within the financial markets industry.
- Own the full flow of challenging data problems starting with data discovery, curation and modelling through to production implementation and monitoring.
- Bring a combination of mathematical, statistical rigor and innovative algorithm design to lead the development of machine learning models, prescriptive analytics and predictive analytics techniques across both structured and unstructured data to build solutions both external and internal customers.
- Remain up-to-date on advanced analytics research and trends, challenge current best thinking, test theories and evaluate feature concepts.
- Develop and model data visualisation components to showcase findings and insights using open-source and commercially available visualisation tools.
- Support the productisation process of data products developed by ASX’s third-party partners.
- Develop validation test cases, detect quality issues, identify their root causes, implement fixes and automate the model development pipeline.
- Provide operational support for the products created through the data platform.

What you've done:

- At least 3 years' experience working as a Data Scientist in a big data environment with expertise in information retrieval, AI, machine learning, natural language processing, data mining and large scale distributed computation.
- Extensive experience in coding using SQL and Python or R.
- Significant demonstrated experience with big data and data analytics technologies such as Hadoop, MapReduce, Hive, Spark, Python, Kafka etc. with a good understanding of Cloudera Big Data Platform.
- Prior experience in statistical and predictive modelling with exposure to supervised and unsupervised machine learning models (e.g. Tensor Flow, clustering, neural networks, decision trees etc.).
- Accredited tertiary qualifications in statistics, maths, computer science or a relevant field.

And if you've got some of this, even better:

- Experience in using Spyder IDE, RStudio, Jupyter, and Cloudera Data Science Workbench.
- Exposure to capital markets datasets and data use cases.
- Operated in a customer-led and commercially-focussed environment.
- Experience in data analysis, modelling and metadata management as well as with model ethics and governance.
- Experience working in a lean agile environment or within a startup.

What you need to enjoy and be good at for this role:

- Solving problems through innovative use of data and technologies with a commitment to quality.
- Working in a fast-paced and iterative environment where outcomes are often unknown at the outset.
- Communicating effectively with technical teams and other stakeholders.
- Being open to taking on new responsibilities which will arise when working in a startup environment.