AI Fundamentals / AI in CRM

Salesforce Certified AI Associate

KNOWLEDGE CHEATSHEET

CHOOSE THE CAPABILITY

- **Predict** — Estimate outcomes
- **Generate** — Create content
- **NLP** — Work with language
- **Vision** — Interpret images

01 AI fundamentals

Learning and models

- **Artificial intelligence (AI)** performs tasks such as recognizing patterns, understanding language and making predictions. A fixed workflow follows predefined rules; machine learning learns patterns from examples.
- **Machine learning (ML)** fits a model to training data. At inference time, the model applies learned patterns to new inputs; learning does not imply human understanding.
- **Neural networks** connect processing units using learned weights and biases. **Deep learning** uses multiple layers; it is a subset of machine learning.
- **Features** are model inputs; a **label** is a known target used in supervised learning. Training learns patterns; evaluation checks performance on separate examples.

Predict, generate and understand

- **Predictive AI** estimates an outcome from existing data: a conversion likelihood, demand forecast or case category. **Classification** predicts categories; **regression** predicts numeric values.
- **Generative AI** creates content such as emails, summaries and article drafts from learned patterns and supplied context. Both predictive and generative AI can use machine learning.
- **Natural language processing (NLP)** works with human language, supporting intent detection, text classification, chat and summarization. **Computer vision** interprets images or video.
- **Hallucinations** are plausible but unsupported generated claims. Check generated content against reliable information before using it in customer communications or decisions.

02 AI capabilities in CRM

Sales productivity and predictions

- **Einstein Lead Scoring** estimates conversion likelihood to help prioritize leads. **Opportunity Scoring** helps assess deal outcomes; scores guide attention and require business judgment.
- **Forecasting** estimates future sales or demand from historical patterns. **Account insights** surface relevant signals; neither a score nor a forecast guarantees an outcome.
- **Call summaries** condense previous conversations and capture key points or follow-up actions. Review the summary before saving it to improve the accuracy of activity records.
- **Einstein Prediction Builder** supports custom predictions from Salesforce data. Define the target outcome and inspect relevant historical records before relying on a prediction.

AI in CRM / Ethical AI

Salesforce Certified AI Associate

KNOWLEDGE CHEATSHEET

APPLY AI ACROSS CRM

- **Sales** — Score and forecast
- **Service** — Classify and assist
- **Marketing** — Personalize
- **People** — Review and act

Service and guided assistance

- **Case Classification** predicts case fields from case information, helping agents categorize and route work. Classification assigns categories; it does not itself resolve every case.
- **Einstein Bots** automate conversational self-service for common requests. NLP can interpret customer intent; provide a clear handoff when the bot cannot resolve the issue.
- **Recommendations** suggest relevant knowledge, products or actions. **Next Best Action** presents contextual recommendations; guided self-service combines suitable recommendations with a workflow.
- **Generative assistance** drafts replies, case summaries and Knowledge content from relevant context. Check facts, tone, permissions and sensitive information before sharing.

Marketing engagement and personalization

- **Personalization** uses relevant customer signals to tailor recommendations and messages. Evaluate satisfaction and engagement alongside consent, fairness and customer control.
- **Einstein engagement capabilities** predict engagement and help optimize communication timing or frequency. Personalized email drafts address content; frequency optimization addresses contact fatigue.
- **Marketing Cloud Engagement models**, such as Einstein Engagement Frequency, exclude demographic and purchased third-party data from their inputs. This safeguard does not eliminate every source of bias.
- **Reporting** describes existing records; predictive scoring and forecasting estimate outcomes. Match the capability to the business goal and measure its effect against a baseline.

03 Ethical considerations of AI

Salesforce Trusted AI principles

- **Responsibility** means protecting human rights and entrusted data while preventing abuse. Assess whether an AI use case should be built, as well as whether it can be built.
- **Accountability** means owning impacts on customers, partners and society. Assign responsible owners, gather feedback and provide a way to report and remedy harm.
- **Transparency** helps people understand AI decisions and limitations. Explain the important factors behind predictions and disclose appropriate information about AI use.
- **Empowerment** pairs AI with human ability and makes it accessible across skill levels, including tools configured with clicks. Train users to make informed decisions.
- **Inclusivity** considers everyone affected. Involve varied perspectives and test with representative data across relevant groups and contexts.

Ethical AI

Salesforce Certified AI Associate

KNOWLEDGE CHEATSHEET

BUILD RESPONSIBLE AI

- **Assess** — Impacts and bias
- → **Protect** — Data and people
- → **Explain** — Decisions and limits
- → **Monitor** — Feedback and harm

Explainability and human oversight

- **Explainability** describes why a model produced an outcome. A **model card** documents intended use, construction, performance and limitations; a prediction explanation addresses an individual result.
- **Human oversight** provides review, correction and escalation. People contribute context and ethical judgment; AI can surface patterns that help people investigate problems.
- **Automation bias** is excessive trust in machine output. Preserve critical thinking, challenge unexpected recommendations and retain human review for consequential decisions.

Recognize bias patterns

- **Bias** can enter through collection, sampling, labels, assumptions or user feedback. An unrepresentative dataset can reproduce unequal treatment even when it contains many records.
- **Association bias** links labels to stereotypes. **Societal bias** reproduces historical prejudice or imposes a narrow set of social values through system design and data.
- **Confirmation bias** reinforces prior assumptions; recommendation feedback loops can repeatedly expose people to similar choices. **Survivorship bias** overlooks people or cases excluded from the observed sample.
- **Interaction bias** arises through people's interactions with a system, including attempts to teach a chatbot harmful language. It differs from simply trusting an automated answer.
- **Sensitive attributes**, such as race or gender, can create discrimination risks. **Proxy variables**, such as location in some contexts, can indirectly encode those attributes.

Mitigation and ethical practice

- **Bias mitigation** combines representative data, careful labels, subgroup evaluation and ongoing monitoring. Removing a sensitive field alone does not remove its proxies or historical effects.
- **Consequence scanning** brings stakeholders together to identify intended and unintended effects, affected groups and safeguards. Turn findings into assigned actions and follow-up checks.
- **Ethical debt** accumulates when known risks are deferred, such as releasing a feature despite harmful bias. Address problems during design and testing instead of relying on later cleanup.
- **Ethical AI practice maturity** progresses from ad hoc work to organized, repeatable processes, then managed, sustainable practice and organization-wide optimization. Leadership, resources and measurement support progress.
- **Monitoring** checks performance and impacts as data and behavior change. Collect user feedback and investigate unequal errors; good overall accuracy can hide poor outcomes for a subgroup.

Ethical AI / Data for AI

Salesforce Certified AI Associate

KNOWLEDGE CHEATSHEET

BUILD RESPONSIBLE AI

Assess Impacts and bias → **Protect** Data and people → **Explain** Decisions and limits → **Monitor** Feedback and harm

Privacy and generative AI safeguards

- **Consent and purpose** govern how customer information is used. Record preferences, respect permitted uses and collect only data needed for the task.
- **Least privilege** limits access to the data and actions each role needs. **Data minimization** limits collection; **retention controls** limit how long information is kept.
- **Accuracy and safety** require checking generated claims and using guardrails against harmful content, bias, toxicity and exposure of personally identifiable information (PII).
- **Honesty** includes respecting data provenance and consent and disclosing autonomously delivered AI content. A confident writing style is not evidence that an output is true.
- **Empowerment and sustainability** favor human control and appropriately sized models. Consider the resources needed to deliver useful results.

04 Data for AI

Data quality dimensions

- **Accuracy** means values reflect reality. **Completeness** means required information is present. A fully populated record can still contain incorrect information.
- **Consistency** means equivalent information uses compatible definitions and formats across records and systems. Standardize product categories, country values and naming conventions.
- **Timeliness** means information is recent enough for the use case. **Reliability** concerns trustworthy sources and dependable collection; old data may poorly represent current customer behavior.
- **Uniqueness** means one intended record per entity. Duplicated contacts can cause repeated outreach, distorted counts and misleading training examples.
- **Validity** means values conform to allowed types, ranges and formats. Valid syntax alone does not prove that an email address or other value is correct.

Data for AI

Salesforce Certified AI Associate

KNOWLEDGE CHEATSHEET

PREPARE DEPENDABLE DATA

- **Define** — Business objective
- → **Assess** — Quality and coverage
- → **Prepare** — Clean and standardize
- → **Maintain** — Owners and monitoring

Assess and prepare the data

- **Start with the business objective** and define the outcome to improve. Then check data availability, relevance, access rights and quality against that use case.
- **Profile the data** with reports and suitable quality tools: count nulls, inspect duplicates, compare category values and examine record age. Establish a baseline before remediation.
- **Clean and standardize** inconsistent values, investigate errors, fill verified gaps and reconcile records across systems. Integrating scattered data requires common definitions, not merely more volume.
- **Training readiness** requires enough relevant, representative examples and dependable labels. Dataset size, freshness and coverage matter together; more poor-quality records do not guarantee better predictions.
- **Training and fine-tuning** both depend on data quality. Keep evaluation examples separate and representative so that measured performance reflects real use.

Salesforce data controls

- **Validation rules** reject records that violate a business condition. To require an email or phone number, flag the case where both fields are blank.
- **Picklists** standardize controlled categories, such as product type or destination country. **Text fields** suit open-ended values, such as a customer's preferred name.
- **Matching rules** identify potential duplicates; **duplicate rules** determine how Salesforce handles them, such as warning or blocking. Review and merge duplicates to clean existing records.
- **Lead data cleanup** includes verifying contact details, updating stale values and filling missing information. Workflow redesign alone does not correct inaccurate records.

Manage quality over time

- **Data management strategy** assigns owners and standards for collection, naming, access, quality, backup, retention and disposal. Monitor the standards after initial cleanup.
- **Quality assessments** establish readiness and reveal gaps before modeling. Reassess as customer behavior, data sources and business objectives change.
- **Poor-quality data** can produce inaccurate predictions, irrelevant recommendations and misleading generated content. It can also amplify bias, waste outreach and damage customer trust.
- **High-quality data** supports more reliable insights and better customer experiences. It reduces avoidable errors but cannot guarantee fairness, eliminate hallucinations or replace human oversight.