Contents

1 Understanding AI and Its Role in Chemistry 1

1.1 Chapter Overview 1
1.2 Core Concepts and Definitions (Chemistry-first) 2
1.2.1 What "AI" Means Here 2
1.3 Four Recurring Task Families 2
1.4 Representations: Giving Chemistry to Models 4
1.4.1 Molecules 4
1.4.2 Reactions 4
1.4.3 Materials 4
1.4.4 Proteins and Complexes 5
1.5 Model Families, Briefly 5
1.6 Validation, Uncertainty, and Applicability Domain 5
1.7 Methods and Workflows (Practical Detail) 6
1.7.1 The Minimum Information Set 6
1.7.2 Property Prediction and Ranking 6
1.7.3 Generative Design Linked to Synthesise-ability 7
1.7.4 Planning Syntheses and Writing Procedures 7
1.7.5 Optimising Reactions and Formulations 7
1.7.6 Documentation and Reproducibility 8
1.8 Case Studies and Recent Signals (2024–2025) 8
1.8.1 Complex Structure Prediction as a Design Accelerant 8
1.8.2 End-to-end Small-molecule Pipelines with Clinical Signals 8

RSC Foundations No. 7
AI Revolution in Chemistry
By Brian McKew

© Brian McKew 2026

Published by the Royal Society of Chemistry, www.rsc.org

1.8.3 Protein Language Models as Programmable Design Aides 9
1.8.4 Materials Discovery at Scale 9
1.8.5 Self-driving Labs and BO Templates 9
1.9 Common Pitfalls and How to Mitigate Them 9
1.9.1 "Great On Paper, Poor in the Lab"—Leakage and Non-causal Shortcuts 9
1.9.2 Out-of-domain Predictions 10
1.9.3 Ignoring Synthesise-ability and Process Constraints 10
1.9.4 Over-trusting a Single Surrogate (*e.g.*, Docking) 10
1.9.5 Inadequate Documentation and Governance 10
1.9.6 Accessibility and Permissions Gaps 10
1.10 Implementation Guidance (Checklists and Decision Frameworks) 10
1.10.1 First Deployment (Prediction/Ranking) 10
1.10.2 Generative Campaign (Design → Make → Learn) 11
1.10.3 Closed-loop Optimisation (Reaction/Formulation) 11
1.11 Key Takeaways 11
References 12

2 Fundamentals of AI-driven Automation in Chemical Research 13

2.1 Chapter Overview 13
2.2 Core Concepts and Definitions 14
2.2.1 Automation *Versus* Autonomy 14
2.2.2 High-throughput Experimentation (HTE) 14
2.2.3 DoE *vs.* BO (and Hybrids) 15
2.2.4 Process Analytical Technology (PAT) and Soft Sensors 16
2.2.5 Digital Twins 16
2.2.6 Self-driving Laboratories (SDLs) 16
2.3 Methods and Workflows 16
2.3.1 Designing an Autonomy-ready Experiment 16
2.3.2 Representing Experiments for Learning 17
2.3.3 Choosing and Configuring Bayesian Optimisation 18

2.3.4 Building a Minimal Self-driving Loop 18
2.3.5 Soft Sensors and Hybrid Control 19
2.3.6 Digital Twins for Planning and Scale-up 20
2.4 Case Studies and Recent Signals (2024–2025) 20
2.4.1 Self-driving Laboratories: From Prototype to Platform 20
2.4.2 Closed-loop Optimisation with Minimal Experiments 20
2.4.3 Joint Condition and Route Discovery 21
2.4.4 Upstream Structure Models Guiding Downstream Autonomy 21
2.4.5 Materials Discovery at Scale 21
2.5 Common Pitfalls and How to Mitigate Them 21
2.5.1 Confounding and Drift 21
2.5.2 Data Leakage and Optimistic Metrics 22
2.5.3 Acquisition Myopia 22
2.5.4 Objective Mis-specification 22
2.5.5 Over-automation Without Safety 22
2.5.6 Under-specified Metadata and Governance 22
2.5.7 Over-promising from Upstream Models 23
2.6 Implementation Guidance (Checklists and Decision Frameworks) 23
2.6.1 Launch Your First HTE Screen 23
2.6.2 Transition from HTE to Closed Loop (BO) 23
2.6.3 Minimal Self-driving Loop Kit 23
2.6.4 RL or MPC Instead of BO? 24
2.6.5 Data Governance Checklist 24
2.7 Key Takeaways 24
References 25

3 Revolutionising Drug Discovery and Development 26

3.1 Chapter Overview 26
3.2 Core Concepts and Definitions 27
3.2.1 From Informatics to AI-native Discovery 27
3.2.2 The Funnel Reframed as Decisions with Uncertainty 27
3.2.3 Data and Representations—What Models See 28
3.2.4 Model Families You Will Actually Use 28

3.3 Methods and Workflows 29
3.3.1 Target Identification and Mechanism Hypotheses 29
3.3.2 Checklist—TID/TV 29
3.3.3 Hit Identification—Model-guided Triage 29
3.3.4 Lead Optimisation—Multi-objective Design in Closed Loop 30
3.3.5 Preclinical Development—ADME/Tox and Developability 32
3.3.6 Clinical Development—Where AI Helps (and Where It Cannot Replace Evidence) 32
3.4 Case Studies and Recent Signals (2024–2025) 33
3.4.1 Case 1—An AI-designed Small Molecule with Phase 2a Signal (2025) 33
3.4.2 Case 2—ML-assisted Antibiotic Discovery Against *A. baumannii* 33
3.4.3 Case 3—Closed-loop Medicinal Chemistry and Reaction Optimisation (2024) 34
3.4.4 Case 4—Structure-guided Design in the AF3 Era 34
3.5 Common Pitfalls and How to Mitigate Them 34
3.5.1 Data Leakage and Unrealistic Validation 34
3.5.2 Out-of-domain Deployment 34
3.5.3 Optimising Proxies, Not Outcomes 34
3.5.4 Synthesis Bottlenecks 35
3.5.5 Black-box Opacity 35
3.5.6 Over-extrapolating Early Clinical Signals 35
3.6 Implementation Guidance (Step-by-step) 35
3.6.1 Launch an AI-enabled Hit-finding Campaign 35
3.6.2 Close the LO Loop (Design → Make → Test → Analyse) 36
3.6.3 Integrate with Preclinical and Clinical Partners 36
3.7 Key Takeaways 36
References 37

4 AI-enabled Materials Discovery and Optimisation 38

4.1 Chapter Overview 38
4.2 Core Concepts and Definitions 39

4.2.1 The Materials Design Problem, Rephrased 39
4.2.2 Representations: What the Model "Sees" 40
4.2.3 Learning Regimes You Will Actually Use 40
4.2.4 Stability, Synthesise-ability, and Process Realism 41
4.3 Methods and Workflows 41
4.3.1 Data Curation and Labelling 41
4.3.2 Feature Engineering and Representation 42
4.3.3 Training, Validation, Uncertainty, Domain 42
4.3.4 Screening, Prioritisation, and Active Learning 43
4.3.5 Generative (Inverse) Design 43
4.3.6 Predicting Synthesise-ability and Processing Windows 44
4.3.7 Closing the Loop with Self-driving Labs (SDLs) 44
4.4 Case Studies and Recent Signals (2024–2025) 44
4.4.1 Scaled Candidate Generation with Rapid Experimental Validation 44
4.4.2 Multi-objective Optimisation in Thin-film Devices 45
4.4.3 Polymer Discovery for Degradability Without Giving Up Strength 45
4.4.4 Catalysts with Performance–Cost–Stability Trade-offs 46
4.5 Common Pitfalls and How to Mitigate Them 46
4.5.1 Shortcut Learning and Data Leakage 46
4.5.2 Label Bias and Fidelity Mismatch 46
4.5.3 Ignoring Synthesise-ability and Processing 46
4.5.4 Over-optimising Single Proxies 46
4.5.5 Uncalibrated Uncertainty and Over-confident Deployment 47
4.5.6 Sustainability as an Afterthought 47
4.5.7 Underspecified Governance 47
4.6 Implementation Guidance (Step-by-step) 47
4.6.1 Stand Up a Composition-to-property Screening Pipeline 47
4.6.2 Add Inverse Design and Synthesis Feasibility 48
4.6.3 Optimise Processing *via* BO (Thin Films or Powders) 48
4.6.4 Embed Sustainability by Design 48
4.6.5 Minimal SDL for Materials (Starter Kit) 48
4.7 Key Takeaways 49
References 49

5 Transforming Chemical Manufacturing with AI 51

5.1 Chapter Overview 51
5.2 Core Concepts and Definitions 52
5.2.1 Where AI Fits in the Plant 52
5.2.2 Essential Vocabulary 52
5.3 Methods and Workflows 53
5.3.1 Data Readiness and Context Building 53
5.3.2 Soft Sensors for CQAs 53
5.3.3 Digital Twins (Empirical and Hybrid) 54
5.3.4 Advisory Optimisation and Control (RTO/MPC/BO) 54
5.3.5 Compliance and Governance 54
5.4 Case Applications 54
5.4.1 Continuous Pharma (Flow Synthesis) 54
5.4.2 Crystallisation with Inline Raman 55
5.4.3 Polymerisation with Soft Sensors 55
5.4.4 Separations (Distillation/Filtration) 55
5.5 Common Pitfalls and How to Mitigate Them 55
5.6 Implementation Guidance (Step-by-step) 56
5.6.1 PAT + Soft Sensor Deployment (Batch or Continuous) 56
5.6.2 Advisory BO for Set-point Tuning (Bounded) 56
5.6.3 MPC with Soft-sensor CQAs 57
5.6.4 Scale-up and Solvent Swap Risk Mitigation 57
5.7 Key Takeaways 57
References 57

6 Enhancing Chemical Safety and Toxicology Through AI 59

6.1 Chapter Overview 59
6.2 Core Concepts and Definitions 60
6.2.1 From QSAR to AI-enabled Safety Science 60
6.2.2 Exposure and Risk, Not Just Hazard 61
6.2.3 Environmental Monitoring and Impact 62
6.3 Methods and Workflows 62
6.3.1 Defendable QSAR Development for Critical Endpoints 62

6.3.2 Read-across with AI Assistance 63
6.3.3 Multi-task and Federated Models 64
6.3.4 Mechanistically Anchored Ensembles 64
6.3.5 Environmental Fate and Effects 65
6.3.6 Monitoring 65
6.4 Case Studies and Recent Signals (2024–2025) 65
6.4.1 Case 1—Deep Learning Lifts Consensus on Tox21 Endpoints 65
6.4.2 Case 2—Skin Sensitisation *via* NAMs: DPRA + KeratinoSens + QSAR 66
6.4.3 Case 3—DILI Composite Models for Risk Triage 66
6.4.4 Case 4—Microplastic Identification and Load Estimation 66
6.5 Common Pitfalls and How to Mitigate Them 67
6.5.1 Data Leakage and Optimistic Validation 67
6.5.2 Out-of-domain (OOD) Deployment 67
6.5.3 Conflating Hazard with Risk 67
6.5.4 Single-surrogate Overreach 67
6.5.5 Uninterpretable Models in Critical Contexts 67
6.5.6 Neglecting Uncertainty 67
6.5.7 Governance Gaps 68
6.5.8 Ethics and Fairness 68
6.6 Implementation Guidance (Step-by-step) 68
6.6.1 Build a Defendable QSAR (Regulatory-relevant) 68
6.6.2 Perform Read-across with AI Support 68
6.6.3 Integrate Hazard and Exposure (Risk-centred) 69
6.6.4 Environmental Impact Workflow 69
6.6.5 Documentation and Governance 69
6.7 Key Takeaways 69
References 70

7 Navigating Ethical Considerations and Regulations 72

7.1 Chapter Overview 72
7.2 Core Concepts and Definitions 74
7.2.1 Bias, Fairness, and Representativeness (Chemistry-specific) 74
7.2.2 Transparency and Explainability 75

7.2.3 Accountability and Human Oversight 75
7.2.4 Safety Envelopes for Autonomy 75
7.2.5 Model Risk and Categorisation 75
7.3 Regulatory Landscape (High-level) 76
7.4 Methods and Workflows 76
7.4.1 Governance-by-design (Lifecycle Template) 76
7.4.2 HAZOP-AI: Hazard Analysis for AI-enabled Chemistry 77
7.4.3 Data Governance and Documentation 77
7.4.4 Uncertainty and Applicability Domain as First-class Controls 78
7.4.5 Security and Privacy in Operational Technology (OT) 78
7.4.6 Supplier and Foundation-model Risk 78
7.5 Case Studies and Recent Signals (2024–2025) 79
7.5.1 Signal 1—EU AI Act Programmes Stand Up AI Management Systems 79
7.5.2 Signal 2—Governing LLM Usage for SOPs and Literature 79
7.5.3 Signal 3—Bias Audits in NAM-first Safety Portfolios 79
7.5.4 Signal 4—Cyber Hygiene for Autonomous Lab Platforms 79
7.5.5 Signal 5—Multi-objective Optimisation with Environmental KPIs 80
7.6 Common Pitfalls and How to Mitigate Them 80
7.6.1 "Great Metrics, Wrong Problem" 80
7.6.2 Opaque Decisions in High-stakes Contexts 80
7.6.3 Over-reliance on Single Surrogates 80
7.6.4 Authority Creep in Autonomy 80
7.6.5 Unbounded LLMs in QMS 81
7.6.6 Model Drift Unnoticed 81
7.6.7 Cybersecurity Gaps at OT/IT Boundary 81
7.6.8 Documentation Debt 81
7.7 Implementation Guidance (Step-by-step) 81
7.7.1 Risk Triage and Governance Plan (Two-hour Workshop) 81
7.7.2 Minimum Documentation Pack (One Model) 82
7.7.3 Validating Models for Chemistry Decisions 82
7.7.4 HAZOP-AI in Practice (One Page per Use-case) 82

7.7.5 Vendor/Foundation-model Diligence 82
7.7.6 Monitoring and Incident Response 83
7.8 Key Takeaways 83
References 83

8 The Future of AI in Chemistry 85
8.1 Chapter Overview 85
8.2 Core Concepts and Definitions 87
8.2.1 Foundation Models for Matter 87
8.2.2 Agentic Systems and Autonomous R&D 87
8.2.3 Cloud-native AI in Chemistry 88
8.2.4 Quantum–Classical Hybrids 88
8.2.5 A Human-centred Future 89
8.3 Methods and Workflows 89
8.3.1 Building a Future-ready Design Stack (Today) 89
8.3.2 Cloud-native Chemistry Pipelines 90
8.3.3 Evaluation Harness for Agentic Tools 90
8.3.4 Hybrid Quantum Experiments (Pragmatic) 90
8.4 Case Studies and Recent Signals (2024–2025) 90
8.4.1 Signal A—Structure Reasoning Becomes "Complex-native" 91
8.4.2 Signal B—Programmatic Protein Design 91
8.4.3 Signal C—Clinical Signal for an AI-designed Small Molecule 91
8.4.4 Signal D—Scaled Materials Discovery with Autonomous Validation 91
8.4.5 Signal E—SDLs as Capabilities 91
8.4.6 Signal F—Cloud R&D Becomes Routine 91
8.5 Common Pitfalls and How to Mitigate Them 92
8.6 Implementation Guidance (Step-by-step) 92
8.6.1 Stand Up a "Future-ready" Project in 90 Days 92
8.6.2 Hiring and Skills 93
8.6.3 Quantum Pilot Discipline 93
8.7 Key Takeaways 93
References 94

9 Practical Implementation Guide 95
9.1 Chapter Overview 95

9.2 Implementation Quick-start: Five Decisions to Unlock in Week One 96
9.2.1 Define the Decision and Action 96
9.2.2 Choose the Validation You Intend to Pass 96
9.2.3 Baseline First 96
9.2.4 Plan Uncertainty and Applicability Domain (AD) 96
9.2.5 Pick the Escalation Path 96
9.3 Data Audit and Preparation (Checklist) 97
9.3.1 Structure and Identity 97
9.3.2 Labels and Protocols 97
9.3.3 Splits and Leakage Controls 97
9.4 Representations: Mapping Problems to Encodings 97
9.4.1 Molecules and Materials 97
9.4.2 Reactions and Procedures 97
9.4.3 Spectra and Images 97
9.4.4 Proteins and Complexes 98
9.5 Model Selection and Benchmarking 98
9.5.1 Baselines 98
9.5.2 Advanced 98
9.5.3 Metrics and Calibration 98
9.6 Validation that Predicts Deployment 98
9.7 When to Use BO and SDL Elements 98
9.7.1 BO 98
9.7.2 SDL (Bounded Autonomy) 99
9.8 Uncertainty and AD in Operations 99
9.9 Troubleshooting: Fastest Routes to Green 99
9.9.1 Leakage and Confounding 99
9.9.2 Hallucinating Assistants 99
9.9.3 Uncalibrated Soft Sensors 99
9.9.4 Objective Mismatch 99
9.10 Playbooks (Step-by-step) 100
9.10.1 Property Prediction (Molecules/Materials) 100
9.10.2 Reaction Optimisation (BO) 100
9.10.3 Materials Screening 100
9.10.4 Safety/Toxicology QSAR 100
9.10.5 Process Soft Sensing 100
9.11 Templates (Insert in Your QMS) 101
9.11.1 Data Card (One Page) 101
9.11.2 Model Card (Two–Four Pages) 101
9.11.3 Operating SOP 101

9.11.4 Monitoring SOP 101
9.11.5 Change-control Form 101
9.12 Key Takeaways 101
References 102

10 Conclusion: Embracing the AI-driven Future of Chemistry 103

10.1 Chapter Overview 103
10.2 Core Concepts and Definitions 105
10.2.1 AI as an Instrument: Decision-centred Chemistry 105
10.2.2 Representations Connect Data to Questions 105
10.2.3 Validation that Predicts Deployment 105
10.2.4 Prospective Confirmation and Value of Experiment 105
10.2.5 Governance that Unlocks Speed 106
10.2.6 Sustainability, Safety and Ethics as Constraints 106
10.2.7 Human–AI Collaboration 106
10.3 Methods and Workflows 106
10.3.1 Method 1—The Decision-first Loop (Repeatable Pattern) 106
10.3.2 Method 2—Education that Compounds (for Individuals and Teams) 107
10.3.3 Method 3—Evaluation Harness for Agentic Tools and Copilots 107
10.3.4 Method 4—Sustainability by Design 108
10.3.5 Method 5—Pragmatic Quantum–Classical Pilots 108
10.4 Key Takeaways 108
References 109

Further Reading 110