

NEO — Artificial Intelligence Policy (Diamond Standard)

17 Aug 2026, 18 min read

Operated by: Nudge Education Ltd · **Version:** v04.26 · **Owner:** Director NEO & Head of School + DSL

NEO BY NUDGE EDUCATION

Artificial Intelligence Policy (Diamond Standard — NEO Implementation) Nudge Education Online

Policy Owner	Director, Nudge Education Online & Head of School (as AI Governance Lead during the initial period), countersigned by the Designated Safeguarding Lead (DSL)
---------------------	---

Approved	April 2026
Review Date	April 2027
Version	04.26
Operating Company	Nudge Education Ltd (Company Number 10192753)
Proprietor	Diego Melo
Accreditation Route	Online Education Accreditation Scheme (OEAS) — accreditation in progress

This policy applies to all learners, staff, practitioners, contractors, volunteers and visitors of Nudge Education Online (NEO). NEO is a fully online alternative provision for learners aged 11–18, operated by Nudge Education Ltd. NEO is not a DfE-registered independent school and is not subject to Independent Schools Inspectorate (ISI) inspection. NEO is pursuing OEAS accreditation only.

1. Why This Matters to NEO

Nudge Education Online (NEO) works with young people who have experienced educational trauma, emotionally based school non-attendance (EBSNA), and systemic exclusion. For these learners, artificial intelligence is not a neutral technology. AI shapes how a young person is seen, measured, and categorised. It can determine what learning paths appear available to them, how their behaviour is interpreted, and what futures they can imagine for themselves. This policy is NEO's implementation of the Nudge Education Diamond Standard for ethical, safe, and rights-led AI use. It applies the Diamond Standard to NEO's specific context as a fully online alternative provision operated by Nudge Education Ltd. It sits alongside NEO's Ten Design Principles for infrastructure, and in particular the three agentic-AI pillars: agentic integrity, consent at every threshold, and collaborative human-machine care. NEO uses this policy to: Protect children's rights — treating children's dignity and developmental freedom as infrastructure, not aspiration. Meet and exceed legal duties — complying with UK GDPR, KCSIE 2026, the ICO Children's Code, the Online Safety Act 2023, and voluntarily applying EU AI Act standards where they better protect children. Prevent harm — recognising that AI can cause harm even when "working as designed". Enable human connection — ensuring technology supports the relationship between learner, named practitioner, and qualified subject-specialist teacher, and never replaces it. Set a sector-leading standard — the Diamond Standard, applied with full operational rigour in NEO. This policy applies to all AI systems used in educational, safeguarding, administrative, or support contexts involving NEO learners. It must be read alongside the NEO Online Safety and Acceptable Use Policy, the NEO Data Protection, Confidentiality and Privacy Policy, the NEO Child Protection and Safeguarding Policy, and the NEO Staff Code of Conduct. NEO is not a DfE-registered independent school and is not subject to ISI inspection. NEO is pursuing OEAS accreditation only. This policy goes beyond the baseline set by OEAS to reflect Nudge Education's Diamond Standard commitment.

2. NEO's Promise

NEO is committed to ensuring that AI serves children, not the other way around. In practice this means: Consent as infrastructure — NEO designs systems where children's consent is structurally preserved, not just procedurally

obtained. Rights before efficiency — Where conflicts arise between institutional convenience and children's rights, children's rights prevail. Human oversight always — No AI system operates without meaningful human review and intervention capacity. No AI agent holds access to NEO systems without a named human authoriser and a time-bound grant (NEO Design Principle 8: agentic integrity). Transparency and legibility — Learners, families, and staff can understand when and how AI is being used. When AI has contributed materially to a document, email, or resource, the contribution is declared. Zero tolerance for prohibited practices — NEO will never use AI for social scoring, manipulation, emotion recognition without explicit safeguards, or any practice that treats children as data to be optimised. Relationship is irreplaceable — The named practitioner (mentor) and qualified subject-specialist teacher remain the primary sources of continuity and judgement in a learner's experience. AI supports those relationships; it never substitutes for them.

3. Legal and Regulatory Framework

3.1 UK Legislation

Children Act 1989 and Children Act 2004 — safeguarding duties. Data Protection Act 2018 and UK GDPR — children's data rights. Data (Use and Access) Act 2025 — effective 19 June 2025, subject to check for any subsequent regulations in force at time of review. Equality Act 2010 — protection from discrimination, including algorithmic discrimination. Online Safety Act 2023 — obligations on user-to-user and search services, content risk assessment. Human Rights Act 1998.

3.2 Statutory Guidance

Keeping Children Safe in Education 2026 (KCSIE), including paragraph 163 (misinformation and disinformation as content risk), paragraphs 165 and 175 (generative AI and the DfE Generative AI Product Safety Expectations, including how filtering and monitoring requirements apply to generative AI), and paragraphs 176–177 (Cyber Security Standards for Schools and Colleges). Working Together to Safeguard Children 2023 (updated May 2025). DfE Product Safety Expectations for generative AI in education (January 2025).

3.3 ICO Children's Code (Age Appropriate Design Code)

All AI systems processing children's personal data must comply with the ICO's legally binding code. NEO applies, as a minimum: Best interests of the child as a guiding principle for every design decision. Privacy by default — no data collection beyond what is necessary for the specified purpose. No profiling or behavioural monitoring without explicit consent and safeguards. No nudge techniques that weaken privacy protections. Data minimisation and purpose limitation, both at system selection and at each use. Parental controls that do not compromise children's privacy rights.

3.4 EU AI Act (Voluntary Compliance for Higher Standards)

NEO treats the EU AI Act as a floor, not a ceiling, and adopts its frameworks voluntarily where they better protect vulnerable children: All child-related AI systems are high-risk by default — requiring human oversight, transparency, and contestability. Eight prohibited practices are banned outright (see Section 6). Transparency obligations — NEO informs learners and parents when AI is used and how it affects them.

3.5 UN Convention on the Rights of the Child (UNCRC)

All AI use at NEO must uphold, at a minimum: Article 3 — best interests of the child. Article 12 — right to be heard. Article 16 — right to privacy. Articles 28–29 — right to education that develops personality and talents. Article 36 — protection from exploitation.

4. Consent Infrastructure

NEO does not treat consent as a checkbox. Consent is infrastructure — a set of structural conditions that must be preserved continuously within any AI system used with learners. Any AI system NEO uses must structurally preserve the following five conditions: Legibility — It must be clear when AI interpretation or inference is occurring. A learner, family member, or member of staff should be able to see that AI is involved, not have to infer it. Refusability — Refusal must be possible without explanation or penalty. Where AI is opt-in, opting out is a first-class choice, not an exception. Reversibility — Consent must be withdrawable with forward effect. Previously processed data is handled in line with retention rules; no further AI processing occurs from the moment of withdrawal. Non-inferability — Silence or presence must not be treated as agreement. Consent is affirmative, specific to a purpose, and documented.

Temporal agency — Learners and families control the pace and depth of AI engagement. A system that escalates processing automatically without fresh consent is not acceptable. If a system cannot support these conditions by design, NEO will not use it. This applies equally to systems NEO procures, systems staff bring to their own workflow, and systems embedded in partner platforms.

5. AI Review Process

5.1 The Three Infrastructure Tests

Before any AI system is used with learners, it must pass three infrastructure tests: Legibility Test — can a member of NEO staff explain to a parent, in plain English, when and how the AI is interpreting their child’s data? If no, the system fails. Refusability Test — can a learner or parent refuse AI processing without negative consequence? If no, the system fails. Reversibility Test — can consent be withdrawn, and all relevant processing stopped, within 24 hours? If no, the system fails.

5.2 AI Review Committee and Risk Assessment

NEO convenes an AI Review Committee comprising the Director / Head of School (as AI Governance Lead during the initial period), the DSL, a senior operational or governance member of staff, and a technical adviser where appropriate. During the initial period, the Committee may be as small as two people (the AI Governance Lead and the DSL); it expands as NEO grows. The minimum is never fewer than two people with a documented conflict-of-interest check. The Committee reviews all proposed AI systems against the three infrastructure tests and the prohibited-uses list (Section 6). High-risk systems require a Data Protection Impact Assessment (DPIA) and a Child Rights Impact Assessment before deployment. Restricted-use systems (including any involving emotion recognition, biometric analysis, or safeguarding-adjacent inference) require written DSL approval in addition to Committee sign-off. Approvals are time-bound. A system approved in January is reviewed at the Committee meeting at least within 12 months; reviews are brought forward if the system changes materially or if a concern arises.

6. Prohibited Uses

NEO will never use AI for: Social scoring — rating or ranking children’s behaviour, compliance, or perceived worth. Manipulation — exploiting vulnerabilities, trauma, developmental stage, or neurodivergence. Emotion recognition — inferring emotion from biometric data (except in narrowly defined, explicitly consented therapeutic contexts, and then only after a DPIA, Child Rights Impact Assessment, and written DSL approval). Biometric categorisation — inferring race, political opinions, trade union membership, religious beliefs, sex life, or sexual orientation. Indiscriminate scraping — harvesting facial images or biometric data from the internet, CCTV, or any other source. Risk prediction — predicting likelihood of offending, re-offending, or other harm based on profiling. Deepfakes — creating or using manipulated audio, images, or video of learners, staff, or families. Subliminal manipulation — techniques designed to materially distort behaviour below the threshold of awareness. Where any AI system is found to engage in a prohibited practice, including as an undisclosed feature of a vendor platform, the system is suspended immediately pending review.

7. Permitted Uses (with Safeguards)

NEO may use AI for: Learning support — personalised learning pathways, accessibility features, content generation — always with qualified human oversight and adaptation before learners encounter the output. Administrative efficiency — scheduling, data entry, report drafting — with no child-level inference that affects decisions about that child. Communication support — translation, text-to-speech, speech-to-text, captioning — with explicit learner and parent consent and clear opt-out. Staff productivity tools — research, lesson planning, documentation — provided no personal data of learners is entered without DPO approval and a completed DPIA. Safer recruitment aids — structured scheduling, note summarisation — never shortlisting or appointment decisions, which remain human. All permitted uses require: Human oversight and review, with a named responsible person. Transparency to learners and parents. Compliance with this policy, the NEO Online Safety and Acceptable Use Policy, and the NEO Data Protection, Confidentiality and Privacy Policy. Regular audit and impact monitoring, at least annually. A declared contribution when AI has materially shaped a document, email, or resource.

8. Safeguarding Integration

AI-related safeguarding concerns are treated as seriously as any other safeguarding issue: Staff must report concerns immediately to the Designated Safeguarding Lead (DSL). Concerns include: AI-generated harmful content, inappropriate AI interactions, privacy breaches, manipulative systems, deepfakes (including those targeting staff or learners), or grooming facilitated by AI. All incidents are logged and reviewed by the AI Review Committee. Where AI systems create safeguarding risks, they are suspended immediately pending review. Safeguarding concerns involving AI are also considered under the NEO Online Safety and Acceptable Use Policy and referred to the police, ICO, or Ofcom under the Online Safety Act 2023 where appropriate.

9. Vendor Management

Any technology partner providing AI systems to NEO must: Sign a Data Processing Agreement (DPA) compliant with UK GDPR. Provide evidence of compliance with this policy and relevant UK and international law. Demonstrate the five conditions of Consent Infrastructure in their system design. Allow audit and inspection of AI systems (including model provenance and training data where commercially viable). Agree to suspend or terminate services if non-compliance is identified. Satisfy NEO Design Principle 8 (agentic integrity) — no unmonitored agent access, no default OAuth grants; every integration is explicitly authorised, logged, and time-bounded. NEO will not use AI systems from vendors who cannot meet these requirements.

10. Null Zones (AI-Free Spaces)

Some interactions are too sensitive for AI processing. NEO establishes Null Zones where AI is completely suspended. The following are Null Zones by default: One-to-one safeguarding disclosures. Therapeutic, counselling, or pastoral wellbeing sessions. Disciplinary, restorative, or other sensitive conversations. Conversations about mental health, self-harm, suicide, bereavement, or trauma. Any interaction where a learner, family member, or member of staff requests AI-free communication. In a Null Zone: session recording auto-summarisation, AI captioning, AI sentiment analysis, and any other AI inference tool are turned off. Where a platform cannot be configured to suspend AI processing, that platform is not used for Null Zone contexts. Staff are trained to recognise when Null Zones are appropriate and how to activate them. The named practitioner is the default authoriser of

Null Zone status for a conversation; the DSL is the default authoriser for safeguarding contexts.

11. Alignment with NEO's Ten Design Principles

This policy is continuous with NEO's infrastructure design. The three agentic-AI pillars (Principles 8, 9, 10) are the operational spine of the Diamond Standard at NEO: Principle 8 — Agentic integrity (no unmonitored agent) — No third-party AI or OAuth tool holds access to NEO systems without explicit, logged, time-bound authorisation by a named human. Third-party OAuth is blocked by default at tenant level. Applies to every AI-related procurement, integration, and staff-bring-your-own-tool decision. Principle 9 — Consent at every threshold — Nothing happens by default. Every authorisation of AI processing — new system, new integration, new purpose, new data set — is a deliberate act with a documented reviewer. Aligns with the Reversibility and Non-inferability conditions above. Principle 10 — Collaborative human-machine care — Automation raises signals; humans make decisions. Enforced technically by not granting autonomous action to AI components, and procedurally by this policy's human-oversight requirement for every permitted use. Where this policy and the NEO Access Control reference architecture differ, the more protective standard applies.

12. Research Lineage — Verse-ality

The Diamond Standard operationalises the protective posture set out in the Verse-ality framework — the agent-safety research authored under The Novacene Ltd. Verse-ality articulates four pillars at the institutional scale: identity non-capture, bounded autonomy, consent as protocol, and agent-to-agent hygiene. The Diamond Standard brings those pillars into kitchen-table-scale everyday practice. This policy translates them into NEO's day-to-day operations. Staff do not need to master Verse-ality to apply this policy. They do need to understand the Diamond Standard and the three infrastructure tests. Links to the underlying research are maintained by the AI Governance Lead and are available on request for staff, commissioners, or researchers.

13. Who This Applies To

Learners at NEO — your rights are at the centre of this policy. Parents and carers — you have the right to know how AI affects your child. Staff, practitioners, qualified teachers, contractors, and volunteers — you must understand and follow this policy. Technology partners and vendors — you must meet the Diamond Standard. Commissioners, local authorities, and partner organisations — you can hold NEO accountable to this standard. This policy is available openly on the NEO website and will be shared with regulators, awarding bodies, and any stakeholder who requests it.

14. Implementation Timeline (From Policy Approval)

14.1 Within 30 Days

Confirm the AI Governance Lead (held by the Director / Head of School during the initial period) and form the AI Review Committee. Inventory all AI systems currently in use at NEO, including any embedded AI features in Google Workspace, Google Classroom, and Google Meet. Communicate this policy to staff, parents, and learners (in age-appropriate versions).

14.2 Within 60 Days

Review all AI systems against this policy and the three infrastructure tests. Suspend any prohibited practices immediately. Complete DPIAs for high-risk systems. Update all vendor contracts that cover AI processing.

14.3 Within 90 Days

Obtain DSL written approval on any restricted-use systems. Establish the AI incident reporting system (integrated with the safeguarding log). Deliver first-cycle staff training on AI, including the Diamond Standard, the infrastructure tests, and prohibited practices. Conduct first learner consultation on AI use.

14.4 Ongoing

AI Review Committee meets at least quarterly once NEO is operational (monthly during the pre-launch and early-launch phase). Quarterly compliance audits covering vendor DPAs, DPIAs, and Null-Zone compliance. Annual policy review and staff training refresh. Annual learner consultation on AI use.

15. Key Definitions

AI system: Software that infers from inputs how to generate outputs (predictions, recommendations, decisions) that influence environments. Aligned with the EU AI Act definition.

Consent Infrastructure: The structural conditions that must be preserved within AI systems to enable genuine, ongoing consent — legibility, refusability, reversibility, non-inferability, and temporal agency.

High-risk AI: AI systems presenting significant risk of harm to health, safety, or fundamental rights. All child-related AI is high-risk by default under this policy.

Null Zone: A conversation, setting, or data context where AI processing is completely suspended to protect particularly sensitive interactions.

Prohibited Practice: One of the eight categories of AI banned under EU AI Act Article 5 (as voluntarily adopted by NEO) due to unacceptable risks to fundamental rights.

AI Governance Lead: The named individual responsible for day-to-day operation of this policy. During the initial period, held by the Director / Head of School.

AI Review Committee: The body responsible for reviewing AI systems before and after deployment. Comprises at least the AI Governance Lead and the DSL.

Agentic-AI pillars: The three NEO Design Principles (8–10) that govern autonomous and semi-autonomous AI in NEO systems: agentic integrity, consent at every threshold, collaborative human-machine care.

16. Living Our Values

This policy is not just about compliance. It is about building a culture where: Children are seen as whole human beings — not data points to be optimised. Human relationship remains irreplaceable — technology supports connection, never replaces it. Rights are infrastructure — embedded in systems, not just stated in policies. Care precedes coherence — learner wellbeing comes before institutional efficiency. Agency belongs to learners — they control how technology shapes their learning and development. The Diamond Standard is not about limiting innovation. It is about civilising it. AI can be a powerful tool for learning, accessibility, and support — but only when it is governed by consent, bounded by rights, and stewarded by humans who understand that no algorithm can replace the care, judgement, and relationship that make education transformative. NEO commits to this standard because the young people NEO serves deserve nothing less.

17. Questions or Concerns

Routes for raising questions or concerns: Staff: contact the AI Governance Lead (Director / Head of School) or the DSL. Parents and carers: contact your child's named practitioner or neo@nudgeeducation.co.uk. Learners: speak to your named practitioner, a parent or carer, or contact Childline on 0800 1111. Partners and commissioners: contact the AI Governance Lead at neo@nudgeeducation.co.uk. All concerns are taken seriously and addressed promptly.

18. Monitoring and Review

AI Review Committee meets at least quarterly (monthly during pre-launch and early operation). This policy is reviewed annually and always when KCSIE, OEAS criteria, the Online Safety Act, UK data protection, or the EU AI Act materially change. An interim review is triggered by any material incident, by any change in the pillar-8/9/10 infrastructure, or by the introduction of a new restricted-use system. Learner, staff, and family voice directly informs revisions. The AI Governance Lead reports to the Proprietor annually on the status of AI systems, incidents, DPIAs, and vendor compliance.

Appendices — To Be Developed

The following appendices are listed as separate NEO documents, to be developed by the AI Governance Lead during the implementation timeline:

- Appendix A — AI Request & Review Form
- Appendix B — Child Rights Impact Assessment Template
- Appendix C — Vendor Due Diligence Checklist
- Appendix D — AI Incident Report Form (integrated with the NEO safeguarding log)
- Appendix E — Learner-Friendly Policy Version (11–14 and 15–18 variants)
- Appendix F — Staff Training Programme
- Appendix G — Data Processing Agreement (DPA) Template
- Appendix H — Glossary and Technical Terms (Full Definitions)
- Appendix I — Detailed Legal and Regulatory Framework
- Appendix J — System-Specific Risk Assessment Examples

Related Policies

This policy should be read alongside: Nudge Education Artificial Intelligence Policy (the Diamond Standard) — the parent organisation-level policy NEO Access Control v1.0 (reference architecture for Principles 8–10) NEO Online Safety and Acceptable Use Policy NEO Data Protection, Confidentiality and

Privacy Policy NEO Child Protection and Safeguarding Policy NEO Staff Code of Conduct NEO Safer Recruitment and Use of Volunteers Policy NEO Behaviour and Regulation Policy NEO Complaints Procedure NEO Digital Consent and AI Safety Parent Guide

Document Control

Version	04.26
Approved	April 2026
Next Review	April 2027
Owner	Director, Nudge Education Online & Head of School
Approver	Proprietor (Diego Melo)
Operating Company	Nudge Education Ltd (Company Number 10192753)

Document control

Field	Value
Version	v04.26
Owner	Director NEO & Head of School + DSL
Status	live
Source file	NEO Policies/NEO - Artificial Intelligence Policy (Diamond Standard – NEO Implementation) v04.26.docx