Please note: Steering Board papers are published to support openness and transparency. These papers may contain draft content and proposals that are subject to discussion, review and approval. Publication does not indicate endorsement, agreement, or adoption as policy.

Applying AI to Open Referral UK

Revision	Date	By	Notes
1	25/09/2026	Paul Davidson, iStandUK	First draft
2	29/09/2026	Paul D	Taking on feedback from within iStandUK
3			External feedback

About **1**
Background **1**
What is AI? **2**
Traditional non-AI approach to deploying Open Referral UK **2**
Opportunities for AI **4**
For a publisher 4
For discovering relevant services 4
For an aggregator 5
For a commissioner 5
AI techniques relevant to ORUK **5**
Governance and safeguards **7**
Principles 7
Sources and further reading **8**

About

This iStandUK discussion paper explores where artificial intelligence (AI) could improve the deployment and impact of Open Referral UK (ORUK).

It considers opportunities across the whole journey, including:

- capturing, classifying and assuring service information
- finding services that fit a person's needs and circumstances

- combining information from multiple publishers
- helping commissioners understand provision and possible gaps

Background

ORUK is a UK profile of the international Open Referral Human Services Data Specification and its associated API. It defines a common structure and exchange mechanism for information about services, organisations, locations and related details.

The Ministry of Housing, Communities and Local Government (MHCLG) invited iStandUK to take on stewardship of ORUK from April 2026, including responsibility for promoting and evolving the standard and supporting tools.

iStandUK's ambition is that councils and other public-sector bodies publish high-quality, assured ORUK feeds that can be combined across a locality, region, nation, topic or audience. The same information can then support many uses, including:

- residents finding services that meet their needs
- link workers and other practitioners identifying suitable support
- commissioners understanding provision and building local capacity and resilience

What is AI?

Artificial intelligence is an umbrella term for computer techniques that perform tasks which normally require aspects of human judgement, such as interpreting language, recognising patterns, making predictions or generating text. AI is not a single product or technique.

For ORUK, it is helpful to distinguish three broad approaches:

- **Rules and conventional software** apply instructions that people have defined. They are suitable for precise checks such as whether a required field is missing or whether a postcode has the expected form. They are usually predictable and easy to explain.
- **Machine-learning and language techniques** learn patterns from examples or represent the meaning of text. They can classify services, identify possible duplicates and find relevant records even when different words are used. Their outputs are usually probabilities or rankings rather than certainties.
- **Generative AI**, including large language models, can interpret and produce natural language. It can support conversations, summarise records and draft suggested changes. It can also produce convincing but incorrect information, so its output must be grounded in source material and checked where accuracy matters.

These approaches can be combined. For example, a Service Finder might use meaning-based search to identify candidate services, apply firm rules for distance or age restrictions, and use generative AI to explain the resulting shortlist in plain language. AI can be an additional layer around reliable ORUK data.

Traditional non-AI approach to deploying Open Referral UK

The guidance on the ORUK website currently describes how an organisation can build, publish and use an ORUK compliant service directory. A typical non-AI approach includes the following stages.

Stage	Established practice	What can go wrong
Capture	A publisher gathers information from service providers, its own research, and proactive outreach. Records may include descriptions, contact details, locations, opening times, costs and eligibility.	Important services or details can be missed. Research and data entry can take considerable time.
Assure and maintain	The publisher checks key details and regularly asks providers to confirm their records.	Requires continuing manual effort. Information may remain wrong or out of date between checks.
Tag	The publisher assigns categories from agreed lists so that services can be grouped, discovered, counted and analysed. Tags may be added manually and/or suggested by matching words in titles and descriptions.	Manual tagging can be expensive and inconsistent. Keyword rules can tag incorrectly: for example, “not suitable during pregnancy” may trigger a pregnancy tag. Different publishers may use classifications that do not align.
Publish	The publisher transforms its directory into the ORUK format and makes it available through an ORUK API. The ORUK Validator can check technical conformance.	Technical conformance does not check that service information is accurate or whether important information is missing, or inconsistent.
Navigate	A website or app asks a person to find a service by ‘drilling down’ through layers of categories.	The person must interpret the categories and decide which services may fit their circumstances. Useful services may have been placed under an unexpected category.

Search	A website or app offers a search form, where a person can search for service names, descriptions, keywords and tags, often with filters for place, category or cost.	A person's wording may differ from the record. Traditional keyword search can return many pages of services that appear to match but are not focused on the person's detailed circumstances, especially when eligibility and locality are buried in descriptions.
Present results	A website or app presents service titles and descriptions, usually as a list or map.	People may need to read many descriptions to understand eligibility, freshness and next steps. Results may not be ordered by likely suitability or explain why each service was returned.
Aggregate	An Aggregator combines information from multiple publishers and tries to identify duplicates and conflicting facts. For example … ● a national service (e.g. the Samaritans Hotline) may be repeated in many local feeds. ● a cross-border service may appear under different identifiers.	Manual deduplication can be expensive and inconsistent.
Analyse provision	Analysts can count services by category, place or provider and compare the results with metrics about local needs.	Missing information, poor tagging, and duplicates can distort comparisons.

Opportunities for AI

For a publisher

Opportunity	**Possible improvement**
Discover candidate services	Monitor digital information sources to identify services that may be candidates for inclusion in a directory, so that the publisher can decide whether and how to create a record.

Assure data and quality	Detect missing or contradictory fields, broken links, unusually old records and inconsistent opening times. Compare a record with information published by the provider, alerting the publisher and drafting questions to the provider for confirmation.
Auto-tag	Suggest terms from one or more controlled lists by interpreting the meaning of a service description. Propose mappings between terminologies used by different publishers.
Extract structured information	Suggest addresses, costs, dates, eligibility criteria, access arrangements and referral routes found in descriptive text.
Geocode	Extract addresses and areas served from prose, suggest coordinates or geographic codes (UPRNs) and highlight conflicts between structured locality fields and descriptions.

For discovering relevant services

Opportunity	**Possible improvement**
Understand circumstances	An AI-assisted Service Finder can have a conversation with an Enquirer or Link Worker to clarify needs, preferences, access requirements, location and realistic travel options. The person should be able to see and correct the resulting summary.
Find candidate services	Meaning-based search can find potentially relevant services even when the person and the service description use different words. For example, “gentle stretching while I'm expecting” may identify an antenatal yoga service.
Check important constraints	Candidate services can be filtered using reliable information about eligibility, locality, cost, access mode, dates and referral route. Constraints found only in descriptive text should be shown as provisional and linked to the source wording.
Explain results	Present a short, ordered list with an explanation of why each service may be relevant.

Improve accessibility	Produce plain-language summaries and support multiple languages in text or speech, while retaining access to the original information and preserving important exclusions and instructions.

For an aggregator

Opportunity	**Possible improvement**
Reconcile many feeds	Compare the contents of multiple feeds to highlight duplicates, flag conflicting facts and help reconcile differing classification systems.
Map between terminologies	Propose relationships between local and shared classifications, distinguishing exact, broader, narrower and related matches.
Coexist with other formats	Incorporate suitable service information from sources such as ● OpenActive ● NHS directories … either by producing an ORUK rendition or by searching across several formats.

For a commissioner

Opportunity	**Possible improvement**
Analyse needs	Confirmed needs described through a Service Finder can be aggregated into common themes and compared with external indicators.
Analyse provision	Analyse services, changes and search results, including searches that return no confirmed match, to identify possible gaps, duplication or changing demand.

AI techniques relevant to ORUK

Technique	**Description**	**Possible ORUK uses**	**Limitations**

Rules and validation	Applies explicit rules that are written by people.	Check required fields, formats, dates and API conformance; enforce known eligibility or geographic rules.	Finds only the conditions anticipated in the rules and cannot judge whether a real-world fact is true.
Natural-language processing and information extraction	Identifies names, addresses, dates, fees, eligibility conditions and other facts in text.	Turn provider descriptions into structured fields; find possible contradictions between fields and prose.	Extracted facts can be incomplete or wrong and should retain a link to the source wording.
Classification	Assigns one or more categories to a record.	Suggest tags from an agreed terminology; map local classifications to shared terms.	A plausible tag may still be incorrect, particularly where wording contains ambiguity.
Embeddings and meaning-based search	Represents the meaning of text so that similar ideas can be found even when they use different words.	Find potentially relevant services; compare descriptions; identify possible duplicates.	Similarity does not establish availability or suitability.
Entity matching	Compares identifiers and other attributes to estimate whether records refer to the same organisation, service or location.	Reconcile records from multiple feeds and highlight possible duplicates.	Similar records can describe genuinely different services; suggested matches require review.
Change and anomaly detection	Highlights unusual values or changes from a previous version or comparable records.	Find stale records, changed provider pages, inconsistent opening times or unexpected gaps.	An unusual record is not necessarily wrong, and a normal-looking record may still be out of date.

Generative AI and large language models	Interprets instructions and produces natural-language responses.	Conduct a service-finding conversation; summarise records; draft provider questions, translations and explanations.	Can invent facts, omit qualifications or vary its answers. Responses must be grounded in current sources.
Geocoding and geographic analysis	Converts addresses into coordinates or areas and compares them with service coverage.	Suggest locations, travel distances and areas served.	use a defined authoritative address and geographic sources. UPRNs from GeoPlace.

Governance and safeguards

The use of AI does not transfer accountability from the organisation operating the service. The AI Playbook for the UK Government recommends safe, responsible and effective use of AI in government. The ICO's guidance on AI and data protection explains how data-protection obligations apply when AI processes personal data. These should be used alongside existing service, information-governance, security, accessibility, procurement and professional controls.

Principles

Safeguard	**What it means for ORUK-related AI**
Clear purpose and accountability	Define the problem, intended users and permitted uses before selecting technology. Name an accountable organisation and owner for each AI-assisted capability, including components supplied by third parties.
ORUK remains the source structure	Keep the ORUK record, its publisher, identifiers, version and update date available. AI-generated text or classifications must not silently replace attributable source information.
Human control over consequential decisions	Use AI to suggest, rank, explain or flag. Do not allow it alone to decide that a person is eligible, refuse support, change a provider's record or merge records where an error could cause harm. Make review proportionate to the risk.

Evidence and uncertainty	Show the source passages, structured fields and dates supporting a suggestion. Distinguish confirmed facts from model inferences and unknowns. Record confidence where it helps reviewers, but do not present a score as proof.
Privacy and data protection	Collect only the personal information needed for the immediate purpose. Define retention and access arrangements, complete a Data Protection Impact Assessment where required, and do not reuse sensitive enquiries for training or analytics without an appropriate lawful basis and safeguards.
Fairness and inclusion	Test whether terminology, ranking, translation or conversational design works differently across places, languages, disabilities and demographic groups. Provide a non-AI route and ensure accessibility is tested with users.
Security	Use approved services and protect records, prompts, logs and credentials. Assess risks such as prompt injection, insecure links, extraction of personal information and unauthorised changes to source data.
Transparency	Tell users when AI is being used, what it contributes and its main limitations. Public bodies should consider publishing an Algorithmic Transparency Recording Standard record where appropriate.
Testing before release	Evaluate against realistic enquiries and records, including ambiguous wording, exclusions, multilingual requests, stale information, duplicates and cases where no service is suitable. Compare performance with the existing process, not with an assumption of perfect data.
Monitoring and correction	Monitor errors, user feedback, data drift, model or supplier changes and unequal outcomes. Provide a clear route to challenge a result, correct a service record and report harm. Keep an audit trail of material suggestions and approvals.
Proportionate procurement and supplier control	Understand where information is processed, whether prompts or records are retained or used for training, how models are changed, and how the organisation can retrieve its data or change supplier. Contractual claims do not replace local testing.

Sources and further reading

- Open Referral UK and developer resources
- Open Referral Human Services Data Specification

- iStandUK and Open Referral UK and the stewardship announcement
- AI Playbook for the UK Government
- Data and AI Ethics Framework
- ICO guidance on AI and data protection
- ICO guidance on explaining decisions made with AI
- Algorithmic Transparency Recording Standard guidance
- NIST AI Risk Management Framework