Natural Resources for Human Health
eISSN: 2583-1194

Original Research

Received 20 July 2026
Revised 28 Aug 2026
Accepted 10 Sep 2026

Natural Resources for Human Health 2026, Vol. 6 (22s): 1691–1704

KEYWORDS:
Generative artificial intelligence; Health information access; Public library; AI anxiety; Trust; AI literacy; Digital divide; Moderated mediation

Generative Artificial Intelligence Services in Public Libraries as a Community Health-Information Gateway: Acceptance, Institutional Trust and the Age Divide

Mi-Ok Jeong[1]*

[1]Department of Library and Information Science, Cheongju University, Cheongju 28503, Republic of Korea
*Corresponding author: miok@cju.ac.kr

ABSTRACT: Public libraries are among the few remaining physical institutions that supply everyday health and life information to an entire community free of charge, and the arrival of generative artificial intelligence (AI) is reshaping how that information is reached. Whether the technology narrows or widens existing information inequalities depends less on the capability of the system than on the disposition of the user. This study examines how acceptance of a library generative AI service translates into expected access to community health information, whether AI anxiety constrains that acceptance, and whether institutional trust buffers the constraint. A cross-sectional survey of 353 public library users in the Republic of Korea measured AI anxiety, AI literacy, facilitating conditions, hedonic motivation, trust, attitude, use intention, recommendation intention and perceived health-information access on five-point Likert scales. Partial least squares structural equation modelling was used to test a serial path from AI anxiety through attitude and use intention to recommendation intention and to perceived health-information access, with trust specified as a moderator and age as a grouping variable. Use intention predicted perceived health-information access ($\beta = 0.448$, $p < 0.001$) and recommendation intention ($\beta = 0.749$, $p < 0.001$). The health-information path was almost twice as strong among users aged 60 years or older as among younger users ($\beta = 0.625$ versus 0.335; $\Delta\beta = 0.290$, $p = 0.002$). AI anxiety was unrelated to attitude ($\beta = 0.080$, $p = 0.281$) and trust did not moderate that relationship ($\beta = 0.001$, $p = 0.972$); trust instead operated as the strongest direct antecedent of attitude ($\beta = 0.444$) and use intention ($\beta = 0.317$). Attitude was predicted by facilitating conditions ($\beta = 0.236$) and AI literacy ($\beta = 0.181$). The findings locate the public health value of library AI at the point where intention converts into expected information access, and they show that this conversion is strongest precisely among the older users whose health-information needs are most acute.

1. INTRODUCTION

1.1 Background

Access to understandable health information is a determinant of health in its own right, not merely a convenience that follows from clinical care [1]. The circulation of unverified health claims during recent public health emergencies demonstrated that the volume of available information and the capacity of citizens to evaluate it have diverged sharply [2]. Instruments developed to measure electronic health literacy consistently show that the ability to locate, appraise and apply health information is unevenly distributed across age, education and occupational groups [3], and that inequality has shifted from whether a person can reach a digital service to how competently that person can use it [4].

Public libraries occupy an unusual position in this landscape. They are publicly funded, physically present in residential neighbourhoods, staffed by information professionals, and free at the point of use. For older adults, caregivers, residents with limited formal education and migrants, a neighbourhood library is frequently the only institution that will help interpret a dense medical document without charge and without a prior appointment. The information a citizen seeks there is rarely confined to bibliographic queries: it concerns a diagnosis received the previous week, the care of an ageing parent, the development of an infant, or a wish to read something consoling during a period of psychological distress.

Generative AI changes the mechanics of that encounter. A library that deploys a conversational agent can answer procedural questions in natural language, recommend reading in response to a sentence describing a feeling rather than a subject heading, summarise a long clinical report, translate foreign-language material, retrieve holdings from a vague description, and direct users to relevant programmes. Professional discussion of these capabilities has moved quickly from speculation to implementation planning [5,6], while recent user surveys report a public that values AI for access and learning yet remains cautious about accuracy, misinformation and privacy [7]. A Korean study of AI chatbot services in public libraries has already shown that acceptance in this setting is driven as much by user disposition as by system capability [10].

What this literature has not yet established is the user-side condition under which such services actually reduce, rather than reproduce, information inequality. Technology acceptance research in library settings has largely transplanted the predictors of the unified theory of acceptance and use of technology, treating performance and effort expectancy as the principal drivers of intention [8,9], and recent library applications follow the same template [10]. Those studies also stop where adoption begins. Intention to use is treated as the terminal outcome, which leaves the question that matters for public health — whether adoption delivers anything — outside the model. Two psychological constructs matter at the margin where hesitation turns into non-use. The first is AI anxiety, a cluster of apprehensions concerning the learning burden of AI tools, the displacement of human work, the social consequences of systems perceived as uncontrollable, and the discomfort of configuring unfamiliar functions [14]. The second is AI literacy, the capacity to recognise how such systems operate, to use them purposefully, to appraise their output critically, and to apply them within ethical limits [16,17].

If anxiety suppresses acceptance, the practical question for a public institution is what it can change. It cannot readily lower a user's dispositional anxiety, and it cannot rapidly raise a population's literacy. It can, however, act on trust. Trust in a service provider is a willingness to accept vulnerability based on the expectation that the provider will behave competently and benevolently [18], and in mediated information environments it is the mechanism that allows users to act despite incomplete understanding of the underlying system [19]. A public library carries an institutional trust endowment that commercial AI platforms do not. Trust is also the construct that recent work identifies as the principal route to public acceptance of AI in government services [20]. Whether that endowment is strong enough to attenuate the effect of anxiety on acceptance is an empirical question with direct operational consequences, and it remains open: studies that model trust in AI services have generally omitted anxiety, and studies of AI anxiety have generally omitted trust [15,23]. Where a moderating role for trust has been specified, it has not been confirmed [21], while anxiety has been shown to moderate the translation of favourable attitudes into motivation [22].

1.2 Theoretical framework and hypothesis development

This study adopts attitude as the proximal mediator between psychological antecedents and behavioural intention, following the decomposed theory of planned behaviour, in which attitude summarises the evaluative consequence of belief structures and transmits them to intention [11,12]. Perceived usefulness and ease of use are retained in spirit through facilitating conditions and literacy rather than reproduced as separate expectancy constructs [13], because the service under study is not yet in routine use and respondents cannot report experience-based expectancies with reliability.

Four exogenous constructs are specified. AI anxiety is expected to depress evaluative judgement of the service, as apprehension about a technology is carried into appraisal of any application of it [14,15,23]. AI literacy is expected to raise that judgement, since users who understand a system's operation and limits can situate its output within realistic expectations [16]. Facilitating conditions, comprising device access, availability of help, institutional support and compatibility with existing tools, are expected to raise attitude by reducing anticipated friction [8,9]. Hedonic motivation, the pleasure anticipated from the interaction itself, is expected to raise attitude in a voluntary leisure-adjacent setting such as a public library [9].

H1. AI anxiety is negatively associated with attitude towards the service.

H2. AI literacy is positively associated with attitude towards the service.

H3. Facilitating conditions are positively associated with attitude towards the service.

H4. Hedonic motivation is positively associated with attitude towards the service.

Attitude is expected to predict use intention, and use intention to predict recommendation intention, since advocacy to others presupposes a personal willingness to use [9,24]. Recommendation intention matters for a public institution because it is the point at which an individual decision becomes a community one: a user who recommends the service to a neighbour extends its reach to people who would not have approached it unprompted.

Use intention is further specified as a predictor of perceived health-information access, the user's expectation that the service will help them locate information about illness and health, obtain information for caregiving and later life, make difficult medical material comprehensible, and find consoling reading during psychological distress. This is the step at which a technology decision inside a library becomes a health-information outcome in the community, and it is the step that library acceptance research, which stops at adoption, has left unmeasured.

H5. Attitude is positively associated with use intention.

H6. Use intention is positively associated with recommendation intention.

H7. Use intention is positively associated with perceived health-information access.

Trust is specified as a moderator rather than an antecedent. The theoretical claim is conditional: anxiety is an expectation of harm or loss of control, and trust is an expectation that the provider will not permit harm, so the behavioural weight of anxiety should depend on how much the user trusts the institution operating the system [18,19]. Where trust is high, the anxious user has an external warrant for proceeding; where it is low, anxiety operates unchecked. The same logic applies, more weakly, to the translation of a favourable attitude into intention.

H8. Trust moderates the relationship between AI anxiety and attitude, such that the negative association weakens as trust increases.

H9. Trust moderates the relationship between attitude and use intention, such that the positive association strengthens as trust increases.

The mediation structure follows from H1, H5 and H6, and the combination of mediation with moderation yields a conditional indirect effect [31].

H10. Attitude mediates the relationship between AI anxiety and use intention.

H11. Attitude and use intention serially mediate the relationship between AI anxiety and recommendation intention.

H12. The serial indirect effect specified in H11 is conditional on the level of trust.

Finally, because the policy concern is distributional, the model is compared across age groups. Older users are, on average, lower in digital operational competence and higher in technology-related apprehension, and they are also the group whose health information needs are most acute [4].

H13. The structural relationships differ between users aged 60 years or older and users younger than 60 years.

The hypothesised model is presented in Figure 1.

Natural Resources for Human Health
eISSN: 2583-1194

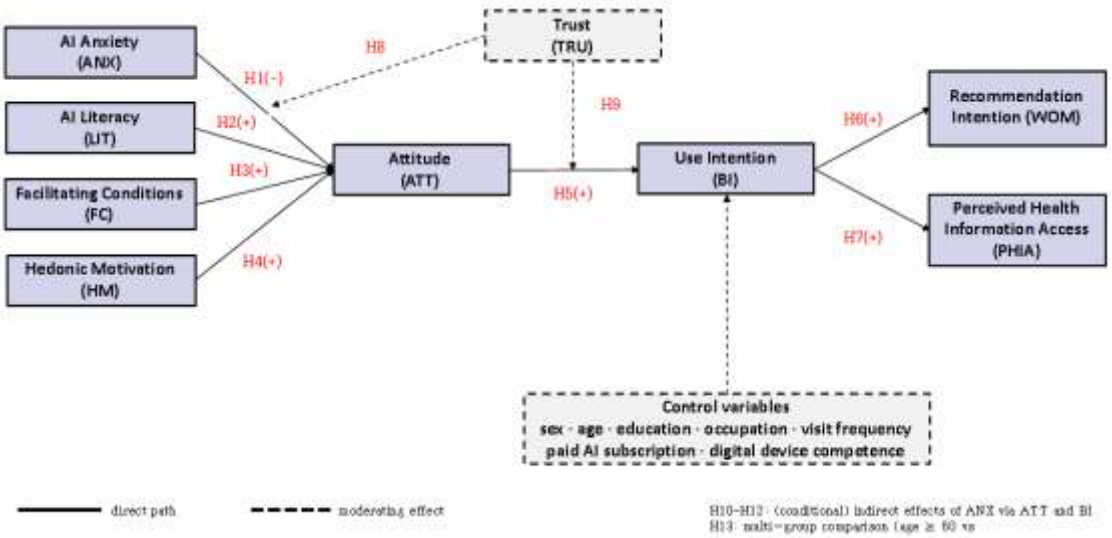

Figure 1. Hypothesised research model. Solid arrows denote direct paths and dashed arrows denote moderating effects. H10 to H12 concern the (conditional) indirect effects of AI anxiety transmitted through attitude and use intention; H13 concerns the comparison of the structural relationships between users aged 60 years or older and users younger than 60 years.

1.3 Aim of the study

This study therefore tests, in a single model, whether AI anxiety constrains the acceptance of generative AI services in public libraries, whether institutional trust buffers that constraint, how acceptance propagates to advocacy and to expected access to community health information, and whether these relationships are equivalent across age groups. The contribution is threefold: it joins two literatures that have developed separately by specifying anxiety and trust in interaction; it extends the dependent variable of library technology acceptance research from adoption to a health-information outcome of public health relevance; and it tests, rather than assumes, that the mechanism operates alike for the older users whose health-information needs are greatest.

2. MATERIALS AND METHODS

2.1 Study Design

A cross-sectional survey design with a scenario-based stimulus was used. Because generative AI services are not yet uniformly deployed across public libraries in the Republic of Korea, respondents could not be assumed to have direct experience of them. All participants therefore read a standardised description of six concrete services before responding to the perceptual items: an AI reference chatbot answering procedural questions in everyday language; AI reading recommendation responding to sentence-form requests, including requests framed in terms of emotional state; AI summarisation and translation of long or foreign-language material; an AI search assistant retrieving holdings from descriptive rather than bibliographic queries; AI learning and assignment support; and AI programme guidance matched to stated interests. The scenario fixes the referent of every subsequent item and holds it constant across respondents, which is a precondition for interpreting between-person variation in attitude and intention.

2.2 Participants and Sampling

The target population comprised users of public libraries in Seoul, Republic of Korea, and the survey was administered face to face at two public libraries in the city. Two screening conditions were applied: use of a public library within the preceding six months, and prior awareness of generative AI tools. Respondents younger than 14 years were excluded. 400 questionnaires were distributed on site and all were collected. Questionnaires with manifestly careless or incomplete responding were set aside in the field, and 365 were entered for analysis. A further 12 cases were removed during screening — 5 for an invariant response across every construct block and 7 for a value outside the response scale on at least one model indicator — leaving 353 cases. No quota was set in advance; the achieved sample contained 124 users aged 60 years or older, which is sufficient for independent

estimation of the measurement and structural models in the smaller group. The author administered the survey personally at both sites; no online channel was used.

2.3 Measurement Instrument

All perceptual constructs were measured on five-point Likert scales anchored at 1 (strongly disagree) and 5 (strongly agree). AI anxiety was measured with four items covering the learning, job-replacement, sociotechnical and configuration dimensions of the original scale [14]; the learning-anxiety item was removed during measurement model purification for the reasons given in Section 3.2. AI literacy was measured with four items covering the awareness, use, evaluation and ethics dimensions [16]. Facilitating conditions were measured with four items on device and environmental readiness, availability of help, institutional support and compatibility, and hedonic motivation with three items on enjoyment, interest and the novelty of the information experience [8,9]. Trust was measured with four items on answer accuracy, information reliability, responsible operation by the library, and benevolence towards the user's interests [18,19]. Attitude was measured with four items [11], use intention with four items [9], and recommendation intention with four items [24].

Perceived health-information access was developed for this study, as no existing scale addresses health-information outcomes of library AI services. Four items were written to cover the principal health-adjacent uses implied by the service scenario: locating health and disease information; obtaining information for caregiving, child-rearing and later life; comprehending difficult health and medical material; and locating consoling material during psychological distress. The items were written by the author from the service scenario and the health-information literature. No expert content-validity panel was convened and no pilot administration was conducted, so the scale enters this study without the prior validation that an established instrument would carry. Because the construct is new and unpiloted, its reliability, convergent validity and discriminant validity are reported in full in Section 3.2 rather than assumed, and its status is revisited among the limitations in Section 3.6.

Control variables were sex, age group, educational attainment, occupation, library visit frequency, paid subscription to a commercial generative AI service, and self-assessed digital device competence. All seven were entered as predictors of use intention; occupation was dummy coded.

2.4 Data Collection Procedure

Data were collected between 15 and 30 September 2026. The author approached each prospective respondent in person at the library and showed them the participation notice printed on the first page of the questionnaire, which set out the purpose of the study, the anonymity of responses, the statutory confidentiality of collected data under Article 33 of the Statistics Act of the Republic of Korea, the voluntary nature of participation, the right to stop at any point without disadvantage, and the researcher's contact details. The questionnaire was handed over only to those who, having been shown and told this, gave their permission to proceed. The instrument required approximately five minutes to complete. Cases were excluded if more than five per cent of the substantive items were missing, if any value on a model indicator fell outside the response scale, or if the respondent gave an invariant response across every construct block. 17 out-of-range entries on control variables were set to missing rather than treated as grounds for exclusion.

2.5 Statistical Analysis

Partial least squares structural equation modelling was conducted with a purpose-written estimator implementing Mode A outer estimation and the path weighting inner scheme, and the procedure was verified against SmartPLS 4 output on the same data. This estimator was selected because the model contains two interaction terms and a serial mediation chain, and because the analytical objective is prediction and explanation of endogenous constructs rather than confirmation of a strict covariance structure [25,26].

The measurement model was evaluated by indicator loadings, with 0.708 as the reference threshold; internal consistency by Cronbach's alpha, the reliability coefficient ρA and composite reliability, with 0.70 as the lower bound for each; convergent validity by average variance extracted, with 0.50 as the lower bound [28]; and discriminant validity by the heterotrait-monotrait ratio of correlations, with 0.85 as the conservative ceiling and 0.90 as the ceiling applicable to conceptually proximal constructs [27]. Common method bias was assessed by a full collinearity test, with variance inflation factors below 3.3 taken as evidence against substantive method variance [29], supplemented by the procedural remedies of item separation and anonymity recommended for self-report designs [30].

Because an interaction term requires the main effect of its moderator to be present, trust was entered as a direct predictor of attitude and of use intention alongside the two interaction terms; these main effects are estimated and reported for completeness but are not hypothesised. Approximate model fit was assessed by the standardised root mean square residual computed on the composite model, with 0.08 taken as the upper bound.

The structural model was evaluated by path coefficients, coefficient of determination, effect size f^2, predictive relevance Q^2 obtained by blindfolding with an omission distance chosen so that it divides neither the number of cases nor the number of indicators, and out-of-sample predictive performance using PLSpredict with ten-fold cross-validation benchmarked against a linear model [33]. Statistical inference used bootstrapping with 10,000 subsamples and bias-corrected and accelerated confidence intervals. Moderation was estimated using the two-stage approach and probed by simple slopes at one standard deviation above and below the mean of the moderator. Conditional indirect effects were estimated at the same moderator values. Because the moderator is specified on two paths of the serial chain, the indirect effect is not a linear function of trust and a single linear index does not apply; the index of moderated mediation was therefore computed as the difference between the conditional serial indirect effects at one standard deviation above and below the mean of trust, and its bootstrap confidence interval was used to test whether the indirect effect varies systematically with trust rather than merely attaining significance in one subgroup [31] [26]. Group comparison by age followed a two-step procedure: measurement invariance of composites was established through configural, compositional and equality-of-means-and-variances assessment [32], after which path differences were tested using the permutation-based PLS multi-group analysis. Reported p values are two-tailed.

3. RESULTS AND DISCUSSION

3.1 Preliminary Analysis

Of the 400 questionnaires distributed and collected, 365 were entered for analysis and 353 were retained after screening, giving a valid response rate of 88.3 per cent of distributed questionnaires. The demographic composition of the retained sample is reported in Table 1. The sample contained 124 respondents aged 60 years or older, which meets the minimum group size required for the multi-group comparison reported in Section 3.5. Construct-level means ranged from 3.67 to 4.12, and no item exhibited skewness or kurtosis beyond the conventional absolute value of 2, indicating distributions suitable for the estimator employed.

Table 1. Demographic profile of the valid sample (N = 353)

Characteristic	Category	n	%
Sex	Male	173	49.0
	Female	172	48.7
Age	14-19 years	6	1.7
	20-29 years	50	14.2
	30-39 years	58	16.4
	40-49 years	65	18.4
	50-59 years	50	14.2
	60-69 years	84	23.8
	70 years or older	40	11.3
Education	High school or below	88	24.9
	Junior college	90	25.5
	University	129	36.5
	Postgraduate	45	12.7
Occupation	Student	49	13.9
	Office or managerial	95	26.9
	Professional	83	23.5
	Self-employed, sales or service	74	21.0

	Production or technical	37	10.5
	Homemaker	2	0.6
	Not employed or retired	13	3.7
	Other	0	0.0
Library visits	Once a week or less	70	19.8
	Twice a week or less	91	25.8
	Three times a week or less	111	31.4
	Four times a week or less	56	15.9
	Five times a week or less	25	7.1
	Daily	0	0.0
Paid AI subscription	Yes	151	42.8
	No	194	55.0
Device competence	Very low	77	21.8
	Low	105	29.7
	Moderate	71	20.1
	High	65	18.4
	Very high	35	9.9

3.2 Measurement Model

The AI anxiety scale required purification. Estimated with all four items, the learning-anxiety item — "learning to use AI tools makes me anxious" — returned a loading of −0.085 on its own construct and reduced Cronbach's alpha to 0.618. Its distribution explains why: 268 of 365 respondents, 73.4 per cent, selected the scale midpoint on it. An item on which three quarters of a sample decline to take a position carries almost no variance to share with its construct. It was removed, and the three remaining items — job replacement, sociotechnical uncontrollability and configuration discomfort — constitute the operational definition of AI anxiety throughout. The narrowing has interpretive consequences that are taken up in Section 3.6.

Indicator loadings, internal consistency and convergent validity for the purified model are reported in Table 2. Loadings ranged from 0.640 to 0.926; three indicators fell below the 0.708 reference value and were retained because composite reliability and average variance extracted exceeded their thresholds for the constructs concerned [25]. Cronbach's alpha ranged from 0.737 to 0.941, ρA from 0.754 to 0.945, composite reliability from 0.833 to 0.957, and average variance extracted from 0.557 to 0.849. The newly developed perceived health-information access scale returned $\alpha = 0.754$, $\rho A = 0.774$, composite reliability 0.836 and average variance extracted 0.561, with its four items loading between 0.743 and 0.752 on a single factor — an even but modest loading pattern consistent with a reflective construct whose indicators are equally peripheral to it. The standardised root mean square residual was 0.120, above the 0.08 criterion; the consequences are addressed in Section 3.6.

Table 2. Descriptive statistics, reliability and convergent validity

Construct	Items	Mean	SD	Loading range	α	ρA	CR	AVE
AI anxiety	3	4.07	0.74	0.763-0.902	0.793	0.795	0.880	0.711
AI literacy	4	3.86	0.75	0.669-0.906	0.847	0.849	0.900	0.694
Facilitating conditions	4	4.03	0.61	0.691-0.803	0.737	0.754	0.833	0.557
Hedonic motivation	3	3.78	0.72	0.728-0.907	0.785	0.827	0.873	0.698
Trust	4	3.77	0.62	0.671-0.868	0.760	0.790	0.845	0.581
Attitude	4	3.91	0.58	0.640-0.876	0.790	0.837	0.859	0.607
Use intention	4	4.01	0.65	0.828-0.857	0.866	0.880	0.908	0.712
Recommendation intention	4	4.12	0.79	0.911-0.926	0.941	0.945	0.957	0.849
Perceived health-information access	4	3.67	0.75	0.743-0.752	0.754	0.774	0.836	0.561

Natural Resources for Human Health
eISSN: 2583-1194

Discriminant validity is reported in Table 3 and is the weakest part of the measurement model. The heterotrait-monotrait ratio between facilitating conditions and hedonic motivation was 0.963, exceeding both the conservative ceiling of 0.85 and the liberal ceiling of 0.90 applicable to proximal constructs [27]: in this sample the two are not empirically distinct, and every statement below that contrasts them is made subject to that qualification. Three further pairs approached the conservative ceiling — trust and attitude (0.848), AI anxiety and AI literacy (0.839), and use intention and recommendation intention (0.818) — of which the last involves two conceptually proximal members of the intention family and satisfies the 0.90 criterion. Perceived health-information access was discriminable from every other construct, its highest ratio being 0.689 with AI literacy. The full collinearity test returned variance inflation factors between 1.74 and 3.64; the maximum exceeds the 3.3 benchmark for excluding common method variance [29] while remaining below the value of 5 at which collinearity is treated as critical, so common method variance cannot be ruled out on this evidence.

Table 3. Discriminant validity (heterotrait-monotrait ratio of correlations)

Construct	1	2	3	4	5	6	7	8	9
1 AI anxiety	—								
2 AI literacy	0.839	—							
3 Facilitating conditions	0.588	0.689	—						
4 Hedonic motivation	0.742	0.786	0.963	—					
5 Trust	0.704	0.654	0.660	0.809	—				
6 Attitude	0.603	0.651	0.665	0.672	0.848	—			
7 Use intention	0.538	0.540	0.635	0.550	0.699	0.714	—		
8 Recommendation intention	0.530	0.456	0.593	0.603	0.661	0.639	0.818	—	
9 Perceived health-information access	0.657	0.689	0.672	0.661	0.506	0.584	0.502	0.456	—

3.3 Structural Model and Hypothesis Testing

Structural estimates are reported in Table 4. The model accounted for 50.1 per cent of the variance in attitude, 52.6 per cent in use intention, 56.1 per cent in recommendation intention and 20.1 per cent in perceived health-information access. Blindfolding returned Q^2 values of 0.278, 0.338, 0.440, 0.094, all above zero, and ten-fold PLSpredict returned Q^2predict values of 0.290, 0.321, 0.449, 0.087. Root mean square errors fell below the linear benchmark for attitude, use intention and perceived health-information access, but not for recommendation intention.

AI anxiety was unrelated to attitude ($\beta = 0.080$, $t = 1.08$, $p = 0.281$), and the point estimate carried the opposite sign to the one hypothesised, so H1 was not supported. AI literacy and facilitating conditions were positively associated with attitude ($\beta = 0.181$, $p = 0.005$; $\beta = 0.236$, $p < 0.001$), supporting H2 and H3. Hedonic motivation was not ($\beta = -0.090$, $p = 0.352$), again with the sign reversed, so H4 was not supported. Attitude predicted use intention ($\beta = 0.466$, $p < 0.001$), supporting H5, and use intention predicted both recommendation intention ($\beta = 0.749$, $p < 0.001$) and perceived health-information access ($\beta = 0.448$, $p < 0.001$), supporting H6 and H7. The effect size of use intention on perceived health-information access was 0.251, the second largest in the model after the path to recommendation intention. Among the exogenous predictors of attitude the largest effect size belonged to facilitating conditions ($f^2 = 0.052$), and all four were small by conventional benchmarks. The unhypothesised main effects of the moderator were the strongest influences on the mediator: trust predicted attitude at $\beta = 0.444$ and use intention at $\beta = 0.317$. None of the seven control variables was related to use intention; the largest standardised coefficient among them was 0.094.

Table 4. Structural model results

Hypothesis	Path	β	SE	t	p	95% CI	f^2	Decision
H1	AI anxiety → Attitude	0.080	0.074	1.08	0.281	[-0.060, 0.224]	0.006	Not supported
H2	AI literacy → Attitude	0.181	0.065	2.80	0.005	[0.056, 0.311]	0.028	Supported

H3	Facilitating conditions → Attitude	0.236	0.059	4.00	< 0.001	[0.129, 0.360]	0.052	Supported
H4	Hedonic motivation → Attitude	-0.090	0.096	0.93	0.352	[-0.280, 0.095]	0.005	Not supported
H5	Attitude → Use intention	0.466	0.067	6.94	< 0.001	[0.329, 0.592]	0.226	Supported
H6	Use intention → Recommendation intention	0.749	0.025	30.56	< 0.001	[0.696, 0.792]	1.278	Supported
H7	Use intention → Perceived health-information access	0.448	0.049	9.17	< 0.001	[0.342, 0.535]	0.251	Supported
H8	ANX*TRU → Attitude	0.001	0.038	0.04	0.972	[-0.071, 0.076]	0.000	Not supported
H9	ATT*TRU → Use intention	0.010	0.045	0.22	0.823	[-0.069, 0.110]	0.000	Not supported

3.4 Moderating Effect of Trust and Moderated Serial Mediation

Neither interaction was significant. The product of AI anxiety and trust on attitude was β = 0.001 (t = 0.04, p = 0.972), so H8 was not supported, and simple slopes were flat across the moderator: the association between AI anxiety and attitude was β = 0.078 at one standard deviation below the mean of trust and β = 0.081 at one standard deviation above it (Figure 2). The product of attitude and trust on use intention was β = 0.010 (p = 0.823), so H9 was not supported.

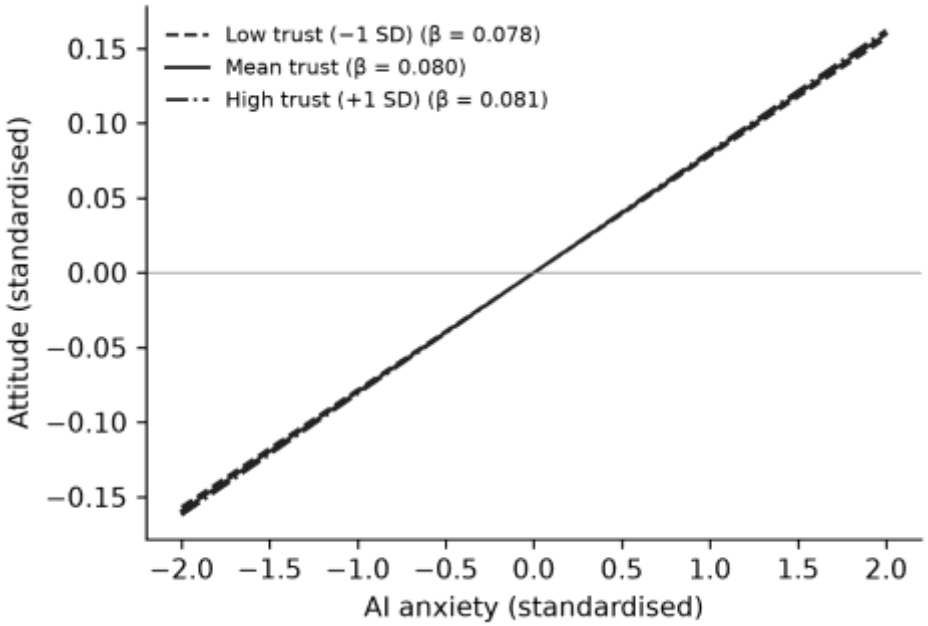

Figure 2. Simple slopes of AI anxiety on attitude at low (−1 SD), mean and high (+1 SD) levels of trust.

Indirect effects are reported in Table 5. Because the first link of the chain was null, nothing was transmitted along it. The indirect effect of AI anxiety on use intention through attitude was 0.037 (95% CI [−0.028, 0.107]) and the serial indirect effect on recommendation intention 0.028 (95% CI [−0.021, 0.081]); both intervals include zero, so H10 and H11 were not supported. The conditional serial indirect effect was 0.027 at low trust and 0.029 at high trust, and the index of moderated mediation was 0.002 (95% CI [−0.050, 0.056]), so H12 was not supported.

Natural Resources for Human Health
eISSN: 2583-1194

Table 5. Mediation and moderated serial mediation

Effect	Estimate	SE	95% CI lower	95% CI upper	Decision
AI anxiety → Attitude → Use intention (H10)	0.037	0.035	-0.028	0.107	Not supported
AI anxiety → Attitude → Use intention → Recommendation intention (H11)	0.028	0.026	-0.021	0.081	Not supported
Conditional serial indirect effect at low trust (-1 SD)	0.027	0.030	-0.030	0.091	Not supported
Conditional serial indirect effect at mean trust	0.028	0.026	-0.021	0.081	Not supported
Conditional serial indirect effect at high trust (+1 SD)	0.029	0.028	-0.024	0.088	Not supported
Index of moderated mediation (high minus low trust) (H12)	0.002	0.027	-0.050	0.056	Not supported

3.5 Multi-Group Analysis by Age

Measurement invariance of composites was examined before group comparison. Configural invariance was established by identical model specification across groups. Compositional invariance held for eight of the nine constructs; for attitude the c value of 0.988 fell marginally below the fifth percentile of its permutation distribution (0.991), and the test of equal composite means and variances indicated partial measurement invariance (Table S1). Partial measurement invariance is sufficient for the comparison of standardised path coefficients [32], but comparisons involving attitude are reported with the caveat that the composite is not demonstrably formed in the same way in both groups.

Permutation-based results are reported in Table 6. Two of the nine paths differed between users aged 60 years or older and younger users, so H13 was supported. The difference of substantive interest concerns the health-information outcome: use intention predicted perceived health-information access at $\beta = 0.625$ among users aged 60 or older and at $\beta = 0.335$ among younger users ($\Delta\beta = 0.290$, $p = 0.002$). AI literacy predicted attitude among younger users ($\beta = 0.297$) but not among older users ($\beta = -0.019$; $\Delta\beta = -0.316$, $p = 0.035$). The remaining seven paths, including the anxiety-to-attitude path and both interaction terms, were statistically equivalent.

Table 6. Multi-group analysis by age group

Path	β (Aged 60 years or older, n = 124)	β (Under 60 years, n = 229)	Δβ	p	Decision
AI anxiety → Attitude	0.104	0.039	0.065	0.696	Equivalent
AI literacy → Attitude	-0.019	0.297	-0.316	0.035	Different
Facilitating conditions → Attitude	0.260	0.219	0.041	0.757	Equivalent
Hedonic motivation → Attitude	-0.067	-0.105	0.038	0.858	Equivalent
Attitude → Use intention	0.352	0.559	-0.207	0.153	Equivalent
Use intention → Recommendation intention	0.811	0.711	0.101	0.064	Equivalent
Use intention → Perceived health-information access	0.625	0.335	0.290	0.002	Different
ANX*TRU → Attitude	0.035	-0.058	0.093	0.261	Equivalent
ATT*TRU → Use intention	-0.098	0.093	-0.190	0.056	Equivalent

3.6 Discussion

The result with the clearest public health consequence is the one the acceptance literature does not usually reach. Users who intend to use a library generative AI service expect it to help them find information about illness and health, obtain information for caregiving and later life, make difficult medical material comprehensible, and locate consoling reading, and the path carrying that expectation was substantial ($\beta = 0.448$, $f^2 = 0.251$). A library's decision to deploy a conversational agent is ordinarily justified in terms of modernisation or convenience. These data allow it to be justified in terms of information access for health, and they quantify how far intention carries into that expectation.

The distributional finding sharpens it. Among users aged 60 years or older the path from intention to expected health-information access was 0.625; among users under 60 it was 0.335, and the difference was the most reliable group contrast in the model ($p = 0.002$). Older users who decide to use the service expect to get markedly more health-information benefit from it than younger users do. This runs against the common presumption that digital services in public institutions serve the already-advantaged first. For the subgroup whose health-information needs are most acute and whose alternatives are fewest, adoption and expected benefit are more tightly coupled, not less. The policy implication is that effort spent bringing older users to the threshold of adoption has a larger expected return per user than the same effort spent on younger ones.

The hypothesis that motivated the study was not supported, and its premise did not hold either. Trust did not moderate the effect of AI anxiety on attitude, but more fundamentally AI anxiety had no effect on attitude to moderate: the coefficient was small, positive and indistinguishable from zero, and nothing propagated down the chain, so every mediated effect in Table 5 contains zero. In this sample the buffering mechanism is not absent because trust failed at it; it is absent because the constraint it was supposed to buffer does not operate.

Two features of the anxiety measure account for much of this, and both are visible in the data rather than inferred. The first is the item that had to be dropped: three quarters of respondents placed themselves at the midpoint of "learning to use AI tools makes me anxious", which is the one item of the four that refers to the respondent's own encounter with the technology. The second is what the three surviving items share — job replacement, uncontrollable social consequences, and discomfort at handling unfamiliar functions. Respondents endorsed them strongly, with a construct mean of 4.07 on a five-point scale, while simultaneously holding favourable attitudes towards the library service (mean 3.91) and reporting high intention to use it (mean 4.01). A person can believe that artificial intelligence threatens employment and still expect a library chatbot to be useful; these are separable judgements about different objects. What the scale captured here is diffuse concern about AI as a social phenomenon, not apprehension about using a particular service. Acceptance research in public institutions should therefore not treat a general AI anxiety scale as a proxy for service-specific use anxiety, and should measure the latter directly if that is the construct of interest.

The null moderation is consistent with the small body of work that has tested it. A 2026 study specifying trust as a moderator of the attitude-to-use relationship in AI adoption likewise failed to support it while finding trust important in importance-performance terms [21], and research on public services positions trust as a direct driver of acceptance rather than as a conditioner of other predictors [20]. The present data point the same way, and emphatically: trust was the strongest single influence on the mediator, predicting attitude at 0.444 and use intention at 0.317 — larger than any hypothesised path into attitude — while contributing nothing as a moderator. For a public library this is the more useful of the two findings. Institutional standing does not neutralise apprehension, but it raises evaluation and intention directly and substantially, and it is a resource the institution already holds.

What else moved attitude was the ordinary machinery of service provision: facilitating conditions carried the largest hypothesised effect, followed by AI literacy, while hedonic motivation carried none. A library AI service appears to be evaluated on whether it looks workable and whether the user feels competent to work it, rather than on whether the encounter promises to be enjoyable. Two qualifications restrain that reading. The effect sizes were all small (f^2 between 0.005 and 0.052), and facilitating conditions and hedonic motivation could not be told apart empirically, so the contrast between them should not be pressed. The age comparison adds a third consideration: AI literacy predicted attitude only among users under 60, its coefficient among older users being indistinguishable from zero. Competence does not convert into favourable evaluation in the older group the way it does in the younger one, which suggests that literacy programmes designed for a general audience will not produce equal returns across the age range.

Several limitations qualify these conclusions, and three constrain them rather than merely annotate them. First, the standardised root mean square residual was 0.120, well above the 0.08 criterion, so the composite model reproduces the observed correlations only approximately and every structural estimate should be read as provisional. Second, facilitating conditions and hedonic motivation returned a heterotrait-monotrait ratio of 0.963: they are not discriminable in this sample, and any claim that one matters and the other does not is unsafe. Third, AI anxiety rests on three items after purification and, as argued above, measures diffuse concern rather than service-specific anxiety; the null result for H1 is a finding about this operationalisation, not a general claim that anxiety is irrelevant to library AI acceptance. Beyond these, several narrower qualifications apply. Compositional invariance was not established for attitude, so group comparisons involving it are tentative, and the maximum full-collinearity variance inflation factor of 3.64 exceeded the 3.3 benchmark for excluding common method variance. Out-of-sample prediction did not beat a linear benchmark for recommendation intention. The perceived health-information access scale is new, explains only 20.1 per cent of its own variance in this model and shows an unusually narrow loading range; it was also developed without an expert content-validity panel or a pilot administration. It should therefore be treated as provisional, and replication with an independently validated instrument is the first thing a reader should ask of this result. The design is cross-sectional, so the mediation sequence is a theoretical ordering rather than a demonstrated temporal one, and responses concern a described service rather than a used one. The sample was drawn face to face from the users of two public libraries in a single metropolitan area, so it represents neither the national population of library users nor those who visit too infrequently to have been present during the fieldwork period. Most importantly for the central claim, perceived health-information access is an expectation, not a measured gain in health-information comprehension. Establishing that the expectation is met requires a longitudinal design in which service use is observed and comprehension is tested directly.

4. CONCLUSION

This study asked whether institutional trust buffers the effect of AI anxiety on the acceptance of generative AI services in public libraries, and what acceptance delivers once it occurs. Trust does not buffer, and in this sample the question's premise did not hold: anxiety about artificial intelligence, as users expressed it, did not depress their evaluation of the library service. Users who worried that AI would take jobs and escape social control nonetheless thought a library AI service was a good idea and intended to use it. What shaped that evaluation was trust in the library, whether the service looked workable, and whether users felt able to work it.

The link with public value is the one worth carrying forward. Acceptance converted strongly into the expectation that the service would make health and everyday-life information easier to reach, and it did so almost twice as strongly among users aged 60 years or older as among younger users. A library's decision to deploy generative AI can therefore be argued not only as modernisation but as health-information policy, with its largest expected return among the users who have the fewest alternatives. For practice the priorities are unglamorous: make the service demonstrably usable, provide help at the point of failure, lean on the institutional standing the library already holds rather than on campaigns to allay general fears about artificial intelligence, and design competence-building for older users specifically rather than scaling down a general programme. Future work should measure service-specific use anxiety separately from diffuse concern about AI, replicate the perceived health-information access scale in an independent sample, and establish longitudinally whether the expectation of better health-information access is borne out in measured comprehension.

ACKNOWLEDGEMENTS

This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors. The author thanks the two participating public libraries in Seoul for permitting the survey on their premises, and the library users who gave their time to answer it.

AUTHOR CONTRIBUTIONS

Mi-Ok Jung: Conceptualization, Methodology, Investigation, Formal Analysis, Writing—Original Draft, Writing—Review & Editing, Supervision.

Mi-Ok Jung is the sole author of this work and conducted every stage of it, including the administration of the survey. There are no co-authors.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

ETHICS STATEMENT

The study was conducted in accordance with the Declaration of Helsinki. It was not submitted to an institutional review board, and no formal determination of exemption was sought or issued; the survey was conducted by the author as an anonymous, non-interventional questionnaire study administered in person.

Consent was obtained face to face at the point of distribution. The first page of the questionnaire carried a written participation notice. It set out the purpose of the study, the anonymity of responses, and the statutory confidentiality of the collected data under Article 33 of the Statistics Act of the Republic of Korea. It further stated that the data would be used only for statistical processing in academic research, that no individual response would be disclosed, that participation was voluntary and could be stopped at any point without disadvantage, and that the questionnaire would take about five minutes. It gave the name, affiliation, electronic mail address and telephone number of the researcher, and closed with a declaration of consent. The author presented this page to each prospective respondent at the library, explained its contents, and handed over the questionnaire only after the person had given permission to proceed. The same page restricted participation to persons aged 14 years or older.

The survey collected no name, contact detail or other directly identifying information at any stage, and no identifier linking a completed questionnaire to an individual was created or retained. Respondents were free to decline without giving a reason and without any consequence for their use of the library.

DATA AVAILABILITY STATEMENT

The data supporting the findings of this study are available from the corresponding author upon reasonable request. The analysis script that reproduces every reported estimate from the raw response file is available from the corresponding author. The data have not been used in, and will not be submitted as part of, any other publication.

REFERENCES

[1] Nutbeam, D., 2000. Health literacy as a public health goal: a challenge for contemporary health education and communication strategies into the 21st century. Health Promotion International 15(3), 259-267.

[2] Shinan-Altman, S., Alon, L., Levkovich, I., 2026. Health information seeking with generative AI among adults with diabetes or prediabetes. DIGITAL HEALTH 12, 20552076261487119. https://doi.org/10.1177/20552076261487119

[3] Norman, C.D., Skinner, H.A., 2006. eHEALS: the eHealth Literacy Scale. Journal of Medical Internet Research 8(4), e27.

[4] Suárez, D., García-Mariñoso, B., 2025. On the verge of a digital divide in the use of generative AI? Telecommunications Policy 49(7), 102997. https://doi.org/10.1016/j.telpol.2025.102997

[5] Ayinde, L., Ebiefung, R., Oladokun, B.D., 2025. Reassessing relevance: public libraries and users' perceptions in the era of generative artificial intelligence. Public Library Quarterly, 1-25. https://doi.org/10.1080/01616846.2025.2610156

[6] Lachaiah, B., Barat, A.H., 2026. Awareness, usage, and perceptions of artificial intelligence among public library users. Public Library Quarterly, 1-31. https://doi.org/10.1080/01616846.2026.2620974

[7] Uzoagba, N.C., Oladokun, B.D., Dumbiri, R., Abdullahi, F., 2026. How public library users perceive the use of generative AI models for information seeking in Nigeria: a cross-sectional survey. Public Library Quarterly, 1-30. https://doi.org/10.1080/01616846.2026.2664362

[8] Venkatesh, V., Morris, M.G., Davis, G.B., Davis, F.D., 2003. User acceptance of information technology: toward a unified view. MIS Quarterly 27(3), 425-478.

[9] Venkatesh, V., Thong, J.Y.L., Xu, X., 2012. Consumer acceptance and use of information technology: extending the unified theory of acceptance and use of technology. MIS Quarterly 36(1), 157-178.

[10] Kang, Y.J., 2026. User acceptance of AI chatbot services in public libraries. The Electronic Library 44(2), 299-318. https://doi.org/10.1108/EL-06-2025-0227

[11] Taylor, S., Todd, P.A., 1995. Understanding information technology usage: a test of competing models. Information Systems Research 6(2), 144-176.

[12] Ajzen, I., 1991. The theory of planned behavior. Organizational Behavior and Human Decision Processes 50(2), 179-211.

[13] Davis, F.D., 1989. Perceived usefulness, perceived ease of use, and user acceptance of information technology. MIS Quarterly 13(3), 319-340.
[14] Wang, Y.Y., Wang, Y.S., 2022. Development and validation of an artificial intelligence anxiety scale: an initial application in predicting motivated learning behavior. Interactive Learning Environments 30(4), 619-634.
[15] Cao, K., Wang, P., Jiang, X., 2026. AI anxiety and adoption intention in higher education based on an extended TAM-UTAUT and PLS-SEM analysis. Scientific Reports 16(1), 3672. https://doi.org/10.1038/s41598-026-35823-9
[16] Wang, B., Rau, P.L.P., Yuan, T., 2023. Measuring user competence in using artificial intelligence: validity and reliability of artificial intelligence literacy scale. Behaviour & Information Technology 42(9), 1324-1337.
[17] Soto-Sanfiel, M.T., Angulo-Brunet, A., Lutz, C., 2025. The scale of artificial intelligence literacy for all (SAIL4ALL): assessing knowledge of artificial intelligence in all adult populations. Humanities and Social Sciences Communications 12(1), 1618. https://doi.org/10.1057/s41599-025-05978-3
[18] Mayer, R.C., Davis, J.H., Schoorman, F.D., 1995. An integrative model of organizational trust. Academy of Management Review 20(3), 709-734.
[19] McKnight, D.H., Choudhury, V., Kacmar, C., 2002. Developing and validating trust measures for e-commerce: an integrative typology. Information Systems Research 13(3), 334-359.
[20] Bunea, O.I., Sabie, O.M., Popescu, R.I., Corboş, R.A., Nastacă, C.C., 2026. Pathways to public trust and acceptance of artificial intelligence in public services: a PLS-SEM and cIPMA approach. Administrative Sciences 16(10), 463. https://doi.org/10.3390/admsci16100463
[21] Xiao, F., Amiri, M., 2026. Extending the technology acceptance model with ethics, trust, and subjective norms: a PLS-SEM analysis of students' AI adoption in higher education. AIP Advances 16(5), 055016. https://doi.org/10.1063/5.0332804
[22] Girgin, Ç., 2026. The moderating role of AI anxiety in the relationship between university students' attitudes toward AI and their motivation. Human Behavior and Emerging Technologies 2026(1), 6243455. https://doi.org/10.1155/hbe2/6243455
[23] Mohamed, M.G., Goktas, P., Khalaf, S.A., Kucukkaya, A., Al-Faouri, I., Seleem, E.A.E.S., Ibraheem, A., Abdelhafez, A.M., Abdullah, S.O., Zaki, H.N., Nashwan, A.J., 2025. Generative artificial intelligence acceptance, anxiety, and behavioral intention in the Middle East: a TAM-based structural equation modelling approach. BMC Nursing 24(1), 703. https://doi.org/10.1186/s12912-025-03436-8
[24] Lu, L., Cai, R., Gursoy, D., 2019. Developing and validating a service robot integration willingness scale. International Journal of Hospitality Management 80, 36-51.
[25] Hair, J.F., Risher, J.J., Sarstedt, M., Ringle, C.M., 2019. When to use and how to report the results of PLS-SEM. European Business Review 31(1), 2-24.
[26] Hair, J.F., Sarstedt, M., Ringle, C.M., Sharma, P.N., Liengaard, B.D., 2024. Going beyond the untold facts in PLS-SEM and moving forward. European Journal of Marketing 58(13), 81-106. https://doi.org/10.1108/EJM-08-2023-0645
[27] Henseler, J., Ringle, C.M., Sarstedt, M., 2015. A new criterion for assessing discriminant validity in variance-based structural equation modeling. Journal of the Academy of Marketing Science 43(1), 115-135.
[28] Fornell, C., Larcker, D.F., 1981. Evaluating structural equation models with unobservable variables and measurement error. Journal of Marketing Research 18(1), 39-50.
[29] Kock, N., 2015. Common method bias in PLS-SEM: a full collinearity assessment approach. International Journal of e-Collaboration 11(4), 1-10.
[30] Podsakoff, P.M., MacKenzie, S.B., Lee, J.Y., Podsakoff, N.P., 2003. Common method biases in behavioral research: a critical review of the literature and recommended remedies. Journal of Applied Psychology 88(5), 879-903.
[31] Hayes, A.F., 2015. An index and test of linear moderated mediation. Multivariate Behavioral Research 50(1), 1-22.
[32] Henseler, J., Ringle, C.M., Sarstedt, M., 2016. Testing measurement invariance of composites using partial least squares. International Marketing Review 33(3), 405-431.
[33] Shmueli, G., Sarstedt, M., Hair, J.F., Cheah, J.H., Ting, H., Vaithilingam, S., Ringle, C.M., 2019. Predictive model assessment in PLS-SEM: guidelines for using PLSpredict. European Journal of Marketing 53(11), 2322-2347.
[34] https://nrfhh.com/index.php/journal/index