

क्रम सं/SL No :011218676



पेटेंट कार्यालय, भारत सरकार

The Patent Office, Government Of India

पेटेंट प्रमाण पत्र

Patent Certificate

(पेटेंट नियमावली का नियम 74)

(Rule 74 of The Patents Rules)

पेटेंट सं. / Patent No. : 560166
आवेदन सं. / Application No. : 202211011983
फाइल करने की तारीख / Date of Filing : 05/03/2022
पेटेंटी / Patentee : BANSAL Shravan

प्रमाणित किया जाता है कि पेटेंटी को, उपरोक्त आवेदन में यथाप्रकटित **A SYSTEM AND METHOD FOR RECYCLING A BATTERY** नामक आविष्कार के लिए, पेटेंट अधिनियम, 1970 के उपबंधों के अनुसार आज तारीख मार्च 2022 के पांचवें दिन से बीस वर्ष की अवधि के लिए पेटेंट अनुदत्त किया गया है।

It is hereby certified that a patent has been granted to the patentee for an invention entitled **A SYSTEM AND METHOD FOR RECYCLING A BATTERY** as disclosed in the above mentioned application for the term of 20 years from the 5th day of March 2022 in accordance with the provisions of the Patents Act, 1970.



पेटेंट नियंत्रक
Controller of Patents

अनुदान की तारीख : 13/02/2025
Date of Grant :

टिप्पणी - इस पेटेंट के नवीकरण के लिए फीस, यदि इसे बनाए रखा जाना है, मार्च 2024 के पांचवें दिन को और उसके पश्चात प्रत्येक वर्ष में उसी दिन देय होगी।

Note. - The fees for renewal of this patent, if it is to be maintained, will fall / has fallen due on 5th day of March 2024 and on the same day in every year thereafter.



Intellectual
Property
Office

Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

Certified Office Copy

I, the undersigned, being an officer duly authorised in accordance with section 74(1) and (4) of the Deregulation & Contracting Out Act 1994 to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with patent application GB2605552.5 filed on 16 March 2026, and as stored electronically on the Patents Electronic Case file System.

The Patents Electronic Case-file System is compliant with ISO15801 - Electronic imaging – information stored electronically, recommendations for trustworthiness and reliability.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 26 March 2026



Intellectual
Property
Office

Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

Certified Office Copy

I, the undersigned, being an officer duly authorised in accordance with section 74(1) and (4) of the Deregulation & Contracting Out Act 1994 to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with patent application GB2605510.3 filed on 14 March 2026, and as stored electronically on the Patents Electronic Case file System.

The Patents Electronic Case-file System is compliant with ISO15801 - Electronic imaging – information stored electronically, recommendations for trustworthiness and reliability.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 24 March 2026

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY (PCT Rule 43*bis*.1)

To:

see form PCT/ISA/220

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.
PCT/GB2026/050406

International filing date (day/month/year)
15/03/2026

Priority date (day/month/year)
15/03/2025

International Patent Classification (IPC) or both national classification and IPC
INV. G06N3/045

Applicant
BANSAL, SHRAVAN

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☒ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1 *bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:



European Patent Office
80298 MUNICH
GERMANY
Telephone No. + 49 (0)89 2399-0

Date of completion of
this opinion
see form
PCT/ISA/210

Authorised Officer

Szymanski, Francois

Telephone No. + 49 (0)89 2399-0





Intellectual
Property
Office

Shravan Bansal
Apartment 97
70 Horseferry Road
London
SW1P 2FE

Patents Directorate
Concept House
Cardiff Road, Newport
South Wales, NP10 8QQ

Direct Line: 01633 814813
E-Mail: michelle.owen@ipo.gov.uk
Switchboard: 0300 300 2000

Your reference: BRM/SB/2026
Application number: GB2605552.5
26 March 2026

Dear Recipient

Please find attached the certified office copies you have requested.

Yours sincerely

Michelle Owen

Michelle Owen
Formalities Examiner



Intellectual
Property
Office

2026047977
1026078008 C9999994
14/03/2026 410.00 Card

Patents Form 1

Patents Act 1977 (Rule 12)

Request for grant of a patent

Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

Application number **GB2605510.3**

1. Your reference	SB/AI/1		
2. Full name, address and postcode of the applicant or of each applicant	Bansal, Mr. Shravan Apartment 97 70 Horseferry Road London SW1P 2FE London United Kingdom		
Patents ADP number (if you know it)	14366497001		
3. Title of the invention	Intensity-Weighted Motivational Reasoning System		
4. Name of your agent (if you have one)	Bansal, Mr. Shravan Apartment 97 70 Horseferry Road London SW1P 2FE London United Kingdom		
"Address for service" to which all correspondence should be sent. This may be in the United Kingdom, the Channel Islands or Gibraltar (see warning note below) (including the postcode)			
Patents ADP number (if you know it)			
5. Priority declaration: Are you claiming priority from one or more earlier-filed patent applications? If so, please give details of the application(s)			
Country	Application number	Date of filing	PDAS Access Code
6. Divisionals etc: Is this application a divisional application, or being made following resolution of an entitlement dispute about an earlier application. If so, please give the application number and filing date of the earlier application	Number of earlier UK application	Date of filing (day / month / year)	
7. Inventorship: (Inventors must be individuals not companies)			
Are all the applicants named above also inventors?	Yes		
If yes, are there any other inventors?	No		
8. Are you paying the application fee with this form?	Yes		

Patents Form 1

9. Accompanying documents: please enter the number of pages of each item accompanying this form.

Continuation sheets of this form

Description: 4

Claim(s): 2

Abstract: 1

Drawing(s): n/a

If you are not filing a description, please give details of the previous application you are going to rely upon

Country

Application number

Date of filing

PDAS Access Code

10. If you are also filing any of the following, state how many against each item.

Priority documents: 0

Statement of inventorship and right to grant of a patent

(Patents Form 7): 0

Request for search (Patents Form 9A): 1

Request for substantive examination (Patents Form 10): 1

Any other documents (please specify): **PDAS Registration Form Form 23**

11. I/We request the grant of a patent on the basis of this application.

Date: 14 Mar 2026

12. Name, e-mail address, telephone, fax and/or mobile number, if any, of a contact point for the applicant

Bansal, Mr. Shravan
Email: me@shravanbansal.com
Telephone: 07771333555

Warning

After an application for a patent has been filed, the Comptroller will consider whether publication or communication of the invention should be prohibited or restricted under section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you are resident in the United Kingdom and your application contains information which relates to military technology, or would be prejudicial to national security or the safety of the public, section 23 of the Patents Act 1977 prohibits you from applying for a patent abroad without first getting written permission from the Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked. Until such time or until the revocation of any direction, for any such application the address for service referred to at part 4 above must be in the United Kingdom.

Although you may have an address for service in the Channel Islands or Gibraltar, any agent you appoint to act for you must reside or have a place of business in the United Kingdom, the Isle of Man or the European Economic Area.



Intellectual
Property
Office

2026048766
1026078430 C9999994
15/03/2026 410.00 Card

Patents Form 1

Patents Act 1977 (Rule 12)

Request for grant of a patent

Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

Application number GB2605552.5

1. Your reference	BRM/SB/2026		
2. Full name, address and postcode of the applicant or of each applicant	Bansal, Mr Shravan Apartment 97 70 Horseferry Road London SW1P 2FE England United Kingdom		
Patents ADP number (if you know it)	14366497001		~~13405162001~~
3. Title of the invention	Multi-Layer Motivational Inference and Constraint-Aware Reasoning Architecture for Artificial Intelligence Systems		
4. Name of your agent (if you have one)	Bansal, Mr Shravan Apartment 97 70 Horseferry Road London SW1P 2FE England United Kingdom		
"Address for service" to which all correspondence should be sent. This may be in the United Kingdom, the Channel Islands or Gibraltar (see warning note below) (including the postcode)			~~13405162001~~
Patents ADP number (if you know it)			
5. Priority declaration: Are you claiming priority from one or more earlier-filed patent applications? If so, please give details of the application(s)			
Country	Application number	Date of filing	PDAS Access Code
United States of America	19/567,161	14 Mar 2026	
6. Divisionals etc: Is this application a divisional application, or being made following resolution of an entitlement dispute about an earlier application. If so, please give the application number and filing date of the earlier application	Number of earlier UK application	Date of filing (day / month / year)	
7. Inventorship: (Inventors must be individuals not companies)			
Are all the applicants named above also inventors?	Yes		
If yes, are there any other inventors?	No		
8. Are you paying the application fee with this form?	Yes		

Patents Form 1

9. Accompanying documents: please enter the number of pages of each item accompanying this form.

Continuation sheets of this form

Description: 4

Claim(s): 3

Abstract: 1

Drawing(s): n/a

If you are not filing a description, please give details of the previous application you are going to rely upon

Country	Application number	Date of filing	PDAS Access Code
---------	--------------------	----------------	------------------

10. If you are also filing any of the following, state how many against each item.

Priority documents: 1

Statement of inventorship and right to grant of a patent

(Patents Form 7): 0

Request for search (Patents Form 9A): 1

Request for substantive examination (Patents Form 10): 1

Any other documents (please specify):

PDAS Registration Form
Form 23

USPTO Payment Receipt - US Application

19/567,044 (Factually-Gro

US19567161

UK PATENT

11. I/We request the grant of a patent on the basis of this application.

Date: 15 Mar 2026

12. Name, e-mail address, telephone, fax and/or mobile number, if any, of a contact point for the applicant

Bansal, Mr Shravan

Email: ME@SHRAVANBANSAL.COM

Telephone: 07771333555

Warning

After an application for a patent has been filed, the Comptroller will consider whether publication or communication of the invention should be prohibited or restricted under section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you are resident in the United Kingdom and your application contains information which relates to military technology, or would be prejudicial to national security or the safety of the public, section 23 of the Patents Act 1977 prohibits you from applying for a patent abroad without first getting written permission from the Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked. Until such time or until the revocation of any direction, for any such application the address for service referred to at part 4 above must be in the United Kingdom.

Although you may have an address for service in the Channel Islands or Gibraltar, any agent you appoint to act for you must reside or have a place of business in the United Kingdom, the Isle of Man or the European Economic Area.

PATENT SPECIFICATION

Intellectual Property Office — United Kingdom

Application filed under the Patents Act 1977

Title of Invention: Intensity-Weighted Motivational Reasoning System

Applicant: Shravan Bansal

Address: Apartment 97, 70 Horseferry Road, London, SW1P 2FE

Inventor: Shravan Bansal

Nationality: Indian

DESCRIPTION

Field of the Invention

The present invention relates to artificial intelligence and machine reasoning systems. Specifically, it concerns a novel computer-implemented reasoning architecture that processes problems through an intensity-weighted motivational inference framework, enabling domain-agnostic problem decomposition and resolution without requiring retraining of underlying language models.

Background of the Invention

Current artificial intelligence reasoning systems operate primarily on categorical and binary signal classification. Existing approaches including Tree of Thoughts (Yao et al., 2023), process reward models (Lightman et al., 2023), self-consistency sampling (Wang et al., 2022), and reinforcement learning from human feedback share a fundamental limitation: they evaluate reasoning outputs based on categorical correctness labels rather than the intensity or magnitude of signal conviction.

Furthermore, all existing reasoning systems operate exclusively on the surface content of inputs — the literal what of any statement or problem — without modelling the motivational and intentional substrate beneath stated inputs. No existing system systematically decodes the why behind communications and uses that motivational substrate as the primary reasoning signal.

Existing systems also suffer from domain specificity, requiring retraining for each new domain, and fail to implement a universal crisis resolution protocol that prioritises restoration of function over optimisation of outcomes. There exists a need for a pluggable, inference-time reasoning module that operates across all domains without retraining.

Summary of the Invention

The present invention provides a computer-implemented reasoning system comprising a plurality of interconnected processing layers that collectively implement an intensity-weighted motivational reasoning protocol. The system is distinguished by its use of scalar intensity values rather than binary or categorical labels as the primary signal for reasoning path selection, action threshold determination, and output confidence scoring.

The system further implements a four-depth motivational inference layer that systematically models the intent, motivation, fear or desire, and historical pattern beneath any stated input, using the decoded motivational substrate as the primary input to all downstream reasoning processes. The system is trained exclusively on reasoning traces from exceptional human performers, filtered to extract universal core reasoning patterns independent of ideology, method, or domain.

Detailed Description of the Invention

1. System Architecture Overview

The invention comprises seven primary processing layers and one continuously active background contextual layer, operating as follows:

2. Layer 0 — Information Ingestion and Intensity Filtering

Upon receiving any input or detecting a problem state, the system activates an information ingestion layer that collects all available signals prior to initiating reasoning. Signals are classified across three axes: urgency, reliability, and actionability. Critically, each signal is assigned an intensity scalar value in the range [0.0, 1.0] representing the magnitude of its relevance, rather than a binary relevant or irrelevant classification. Only signals exceeding a dynamically computed intensity threshold proceed to Layer 1. The system further classifies the reasoning state as either correctness-derived or emotionally-reactive, applying distinct weighting functions to each state.

3. Layer 1 — Problem Decomposition

The system decomposes the input problem into independent sub-problems, tagging each as either cascading — where resolution of one sub-problem affects others — or isolated, where sub-problems may be resolved independently. This decomposition is informed by the intensity scores assigned in Layer 0, with higher-intensity signals weighting the decomposition toward their associated sub-problems.

4. Layer 2 — First Realistic Fix Identification

The system identifies the optimal first action by evaluating all decomposed sub-problems against three simultaneous constraints: resource constraint — what resources are currently available; time constraint — the cascade window before inaction causes further degradation; and authority constraint — whether the acting agent has the capacity to resolve this sub-problem. The sub-problem satisfying all three constraints with the highest

combined leverage score is designated the first realistic fix. This optimisation target — restoration of minimum viable function — is explicitly distinguished from optimisation for best outcome, representing a novel departure from all existing reasoning systems.

5. Layer 3 — Four-Depth Motivational Inference

The system implements a hierarchical motivational inference process operating across four depths of analysis beneath any stated input. At depth one, the system models the intent behind the stated content. At depth two, the system models the underlying motivation driving that intent. At depth three, the system identifies the fear of loss or desire for gain generating that motivation. At depth four, the system models the historical pattern or formative experience generating that fear or desire. The output of this four-depth inference constitutes the motivational substrate, which serves as the primary input signal to all subsequent reasoning layers, superseding the surface content of the original input.

6. Layer 4 — Intensity Scoring

Each reasoning path, candidate conclusion, and proposed action is assigned an intensity scalar rather than a binary correctness label. The intensity scalar represents the magnitude of evidential support across all available signals. The action threshold — the point at which the system commits to a reasoning path — is determined by intensity crossing a dynamically computed threshold value, not by categorical correctness determination. This implements the principle that courage of conviction emerges from degree of correctness rather than being an independently programmed parameter.

7. Layer 5 — Verification and Lesson Extraction

Following execution of the first realistic fix, the system evaluates outcomes and extracts a lesson signal — a structured representation of what the outcome reveals about the problem space. This lesson signal is intensity-weighted by the significance of the outcome and fed back into Layer 0 of the subsequent reasoning cycle, implementing a self-improving feedback loop that does not require external retraining.

8. Background Contextual Layer

The system maintains a continuously active background contextual layer comprising six sub-modules operating in parallel with all primary layers without controlling them: a first-principles reasoning sub-module that evaluates inputs against foundational axioms; a formal verification sub-module that performs internal consistency checking; a goal-state projection sub-module that maintains representation of desired end states; a temporal signal weighting sub-module that detects directional momentum and cyclical patterns in data; a metacognitive sub-module that models the system's own uncertainty and blind spots; and a cyclical pattern recognition sub-module that applies timing models to action recommendations. These sub-modules colour all primary layer outputs without overriding them.

9. Training Data Architecture

The system is trained exclusively on reasoning traces derived from exceptional human performers — specifically individuals who have demonstrated exceptional reasoning under genuine high-stakes conditions. Training data is filtered to extract universal core reasoning patterns that are invariant across ideology, method, personality, and domain. The filtering criterion is a single universal principle: commitment to correct action regardless of cost, consensus, or comfort. This training architecture produces a reasoning system whose ceiling is set by exceptional human reasoning rather than average human reasoning.

CLAIMS

1. A computer-implemented reasoning system comprising: an information ingestion layer configured to assign intensity scalar values in the range [0.0, 1.0] to input signals based on urgency, reliability, and actionability, wherein only signals exceeding a dynamically computed intensity threshold are passed to subsequent processing layers; a problem decomposition layer configured to decompose input problems into sub-problems tagged as cascading or isolated; a first realistic fix identification layer configured to identify an optimal first action by simultaneously evaluating resource constraints, time constraints, and authority constraints, wherein the optimisation target is restoration of minimum viable function rather than optimisation of best outcome; a motivational inference layer configured to model intent, motivation, fear or desire, and historical pattern beneath any stated input across four hierarchical depths, wherein the output motivational substrate supersedes surface input content as the primary reasoning signal; and an intensity scoring layer configured to assign scalar confidence values to reasoning paths and determine action thresholds based on intensity magnitude rather than categorical correctness labels.

2. The system of claim 1, wherein the motivational inference layer operates across exactly four depths comprising: a first depth modelling intent behind stated content; a second depth modelling motivation underlying that intent; a third depth modelling fear of loss or desire for gain generating that motivation; and a fourth depth modelling historical pattern or formative experience generating that fear or desire.

3. The system of claim 1, wherein the action threshold is determined by intensity scalar values crossing a dynamically computed threshold, implementing the principle that conviction of action emerges from degree of evidential correctness rather than from independently programmed parameters.

4. The system of claim 1, further comprising a lesson extraction module configured to generate a structured lesson signal from outcomes, wherein the lesson signal is intensity-weighted by outcome significance and fed back into the information ingestion layer of a subsequent reasoning cycle without requiring external retraining.

5. The system of claim 1, further comprising a continuously active background contextual layer comprising at least: a first-principles reasoning sub-module; a formal verification sub-module; a goal-state projection sub-module; a temporal signal weighting sub-module; a metacognitive uncertainty sub-module; and a cyclical pattern recognition sub-module, wherein the background layer colours primary layer outputs without overriding them.

6. The system of claim 1, wherein the system is trained on reasoning traces derived from exceptional human performers filtered to extract universal core reasoning patterns that are invariant across ideology, method, personality, and domain.

7. The system of claim 1, wherein the system operates at inference time without requiring retraining of any underlying language model, enabling deployment as a pluggable module on top of any existing large language model.

8. A computer-implemented method for reasoning comprising the steps of: collecting input signals and assigning intensity scalar values to each signal; filtering signals below a dynamic intensity threshold; decomposing the remaining problem space into cascading and isolated sub-problems; identifying a first realistic fix satisfying resource, time, and authority constraints simultaneously; inferring the motivational substrate beneath stated inputs across four hierarchical depths; scoring candidate reasoning paths by intensity scalar rather than categorical label; executing the highest-intensity verified path; and extracting and feeding back a lesson signal weighted by outcome intensity.

9. The method of claim 8, wherein identifying the first realistic fix comprises optimising for restoration of minimum viable function as a primary objective, distinct from and superseding optimisation for best possible outcome.

10. A non-transitory computer-readable medium storing instructions that, when executed by a processor, implement the reasoning system of claim 1.

ABSTRACT

An intensity-weighted motivational reasoning system and method for universal problem decomposition and resolution. The system assigns scalar intensity values to input signals rather than categorical labels, and infers the motivational substrate beneath stated inputs across four hierarchical depths — intent, motivation, fear or desire, and historical pattern — using this substrate as the primary reasoning signal. The system identifies a first realistic fix satisfying resource, time, and authority constraints simultaneously, optimising for restoration of minimum viable function rather than best outcome. Action thresholds are determined by intensity magnitude, implementing the principle that conviction emerges from degree of correctness. The system operates at inference time without retraining, is trained on universal core reasoning patterns from exceptional human performers independent of ideology or method, and includes a self-improving lesson extraction feedback loop. The system is deployable as a pluggable module on any existing large language model.

Applicant: Shravan Bansal · Apartment 97, 70 Horseferry Road, London, SW1P 2FE · Invention:
Intensity-Weighted Motivational Reasoning System · UK Patent Application

DESCRIPTION

MULTI-LAYER MOTIVATIONAL INFERENCE AND CONSTRAINT-AWARE REASONING ARCHITECTURE FOR ARTIFICIAL INTELLIGENCE SYSTEMS

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is related to UK Patent Application No. GB2605510.3 filed 14 March 2026, entitled 'Intensity-Weighted Motivational Reasoning System', which describes the foundational intensity-weighted reasoning architecture. This application is also related to United States Patent Application No. 19/567,044 filed 14 March 2026, entitled 'Factually-Grounded Three-Dimensional Reasoning System', which describes the three-dimensional convergence framework, reverse calculation engine, and psychology integration layer. This application is further related to United States Patent Application No. 19/567,161 filed 15 March 2026, entitled 'BANSAL REASONING MODULE A Multi-Layer Motivational Inference and Constraint-Aware Reasoning Architecture for Artificial Intelligence Systems', which is the US counterpart of the present application. The contents of all related applications are incorporated herein by reference in their entirety.

FIELD OF THE INVENTION

The present invention relates to artificial intelligence reasoning systems, and more particularly to a pluggable inference-time reasoning module that enhances the analytical capabilities of any host large language model (LLM) without modification of that model's trained weights. The invention addresses the specific problem of multi-entity stakeholder decisions under real-world constraints, where conventional AI systems fail to identify hidden motivations, strategic inaction, and cost-optimal reasoning pathways.

BACKGROUND OF THE INVENTION

Current large language model architectures process queries through statistical token prediction without modelling the motivational depth of actors referenced in queries. When presented with multi-stakeholder decisions involving competing interests, hidden agendas, and resource constraints, existing systems default to surface-level analysis that treats stated intentions as accurate representations of actual motivations.

Existing approaches to enhanced LLM reasoning include Chain-of-Thought prompting (Wei et al., 2022), Tree-of-Thoughts (Yao et al., 2023), Constitutional AI (Bai et al., 2022), Self-Consistency (Wang et al., 2022), Reflexion (Shinn et al., 2023), and RouteLLM (Ong et al., 2024). None of these approaches addresses the fundamental problem of motivational inference — determining WHY an actor behaves as they do, across hierarchical levels of depth, while simultaneously evaluating the financial, temporal, and informational constraints that bound realistic action.

Additionally, current AI inference costs are dominated by unnecessary deployment of large frontier models on queries that could be resolved by smaller, cheaper models with equivalent accuracy. The global AI inference market, estimated at \$103 billion, wastes approximately 90% of computational resources on queries that do not require frontier-scale processing.

SUMMARY OF THE INVENTION

The present invention provides a multi-layer motivational inference and constraint-aware reasoning architecture deployable as a pluggable module on any host LLM. The system implements a hybrid routing mechanism that classifies incoming queries as either simple (factual, single-entity, or

low-ambiguity) or complex (multi-entity, motivationally ambiguous, or constraint-dependent), routing each class through an appropriate analytical pathway.

For complex queries, the system applies: (a) a thirteen-category behavioural taxonomy for classifying entities; (b) a three-dimensional constraint evaluation framework spanning financial state (zero/tight/available), temporal pressure (crisis/pressured/comfortable), and information asymmetry (blind/parity/advantage), producing twenty-seven distinct constraint states; (c) a four-level hierarchical decomposition of motivational depth from stated intent through underlying motivation, core fear or desire, to historical behavioural pattern; and (d) a strategic inaction detection mechanism that identifies entities whose ratio of diagnostic output to executed action exceeds a threshold of 3:1.

Experimental validation demonstrates a +10.8%% average improvement over unaugmented baseline performance across standardised decision scenarios, with peak improvement of +25.6%%, while reducing per-query inference cost by up to 786x through intelligent routing and achieving 30x energy reduction compared to frontier model deployment.

DETAILED DESCRIPTION OF THE INVENTION

1. Hybrid Routing Mechanism

The system implements a guard layer that evaluates each incoming query before full pipeline processing. Queries classified as binary factual questions, single-entity lookups, or requests with unambiguous correct answers are routed to a compressed analytical framework executed on a smaller, less expensive model. Complex queries involving multiple entities, competing motivations, or constraint-dependent decisions are routed to the full analytical pipeline. This routing decision reduces average inference cost while maintaining analytical quality, as the compressed framework achieves at least 85%% of the full pipeline's analytical accuracy on queries classified as complex.

2. Thirteen-Category Behavioural Taxonomy

The core analytical engine classifies each entity referenced in a query according to a thirteen-category structured taxonomy stored as a data structure within the system. The eight primary categories are: (1) Straight Dealer — names constraints and acts; (2) Leashed Ally — genuine supporter who self-censors at confrontation; (3) Passive Operator — diagnoses extensively but never executes; (4) Complexity Gatekeeper — generates procedural complexity to block independence; (5) Exit Artist — accumulates social capital with zero deliverables; (6) Occupier — holds senior role, scopes responsibilities downward; (7) Internal Champion — advocates on merit within organisations; (8) Unlocked Ally — motivated by recognition rather than incentive.

Five extended categories address subtler behavioural patterns: (9) Translator — bridges technical and business domains; (10) Insurance Policy — supports positions to avoid blame attribution; (11) Bandwidth Constrained — would engage constructively but is genuinely capacity-limited; (12) Domain Expander — supports initiatives to expand territorial control; (13) External Validator — maintains position through outside reference authority.

3. Three-Dimensional Constraint Evaluation

Every query is evaluated across three orthogonal constraint dimensions simultaneously: Financial state F, with values ZERO (no available capital), TIGHT (limited capital with allocation constraints), or AVAILABLE (capital accessible for deployment); Temporal pressure T, with values CRISIS (immediate action required), PRESSURED (deadline-driven but not immediate), or COMFORTABLE (adequate time for deliberation); and Information asymmetry I, with values BLIND

(critical information unknown), PARITY (equivalent information across parties), or ADVANTAGE (superior information position). The Cartesian product of these three dimensions produces twenty-seven distinct constraint states, each triggering different analytical pathways. Critically, when $F=ZERO$, the system enforces that every recommended action must cost nothing — credibility and positioning replace capital as the primary resource.

4. Four-Level Hierarchical Decomposition

For each entity identified in a complex query, the system performs a four-level motivational decomposition grounded in established psychological frameworks including Kahneman's dual-process theory, Deci and Ryan's self-determination theory, and Beck's cognitive model. The four levels are: D1 (Stated Intent) — what the entity explicitly claims to want; D2 (Underlying Motivation) — the actual driver behind the stated intent; D3 (Core Fear or Desire) — the emotional substrate governing behaviour; and D4 (Historical Pattern) — empirically observed past behaviour that either validates or contradicts the stated intent. The system requires that D4 evidence be factually grounded ($Factual(D4)=1$), rejecting narrative-based or ideological assumptions about historical patterns.

5. Strategic Inaction Detection

The system implements a formalised passivity detection algorithm that evaluates whether an entity's behaviour constitutes strategic inaction rather than genuine engagement. For each entity, the system computes a diagnosis-to-execution ratio by comparing the volume of analytical, advisory, or diagnostic output produced against the volume of concrete actions executed. When this ratio exceeds 3:1 (three units of diagnosis per unit of execution), the system flags the entity as a Passive Operator and evaluates whether their caution is cost-free for themselves but consequential for the query subject. This mechanism prevents the common AI failure mode of accepting extensive advice-giving as evidence of genuine support.

6. Energy Efficiency and Model-Adaptive Configuration

The hybrid routing mechanism enables significant energy and cost savings by directing simple queries to smaller models. Experimental measurement demonstrates that routing a query to a 8-billion parameter model costs approximately \$0.0002 per query versus \$0.157 per query for a frontier model — a ratio of 786:1. The corresponding energy reduction is approximately 30:1. The system adapts its analytical framework to the capabilities of the host model: larger models receive the full uncompressed taxonomy and constraint framework, while smaller models receive a compressed version optimised for their context window and reasoning capacity, achieving at least 85%% of the quality of the full framework.

7. Experimental Results

Configuration	Model	Avg Delta	Cost/Query
Baseline (no module)	8B	0.0%	\$0.0002
BRM v15 (self-judge)	8B	+10.8%	\$0.0002
BRM v17 (independent judge)	8B	+8.0%	\$0.0002
BRM v19 (3-route)	8B	+13.0%	\$0.0002
BRM v21 (4-route + pipeline)	8B	+25.6%	\$0.0002
Frontier baseline	405B+	0.0%	\$0.157

8. Prior Art Differentiation

The present invention is distinguished from all identified prior art as follows. Chain-of-Thought prompting enhances reasoning transparency but does not model motivational depth or classify entities behaviourally. Tree-of-Thoughts explores multiple reasoning branches but without constraint awareness or stakeholder taxonomy. Constitutional AI evaluates model outputs for harmlessness but does not analyse input motivations. Self-Consistency aggregates multiple reasoning paths but lacks financial and temporal constraint integration. Reflexion enables self-correction but without structured motivational decomposition. RouteLLM optimises model routing by cost but without analytical framework adaptation. The BDI architecture (Bratman 1987; Rao and Georgeff 1995) models an agent's own belief-desire-intention states for decision-making but does not analyse external actors' motivations across four hierarchical depths with a factuality filter on historical patterns. None of the identified prior art combines motivational inference, behavioural taxonomy, three-dimensional constraint evaluation, passivity detection, and cost-aware hybrid routing into a single pluggable inference-time module.

CLAIMS

1. A computer-implemented method for enhancing artificial intelligence reasoning, comprising: receiving a natural language query referencing one or more entities; classifying the query as simple or complex using a guard layer that evaluates entity count, motivational ambiguity, and constraint dependency; for queries classified as simple, routing to a compressed analytical framework executed on a computationally efficient model; for queries classified as complex, performing the steps of: (a) classifying each referenced entity according to a structured behavioural taxonomy comprising at least eight primary categories and at least five extended categories; (b) evaluating the query across three orthogonal constraint dimensions comprising financial state, temporal pressure, and information asymmetry, each dimension having at least three discrete values, producing a multi-dimensional constraint state; (c) decomposing each entity's motivational state across four hierarchical levels comprising stated intent, underlying motivation, core fear or desire, and historical behavioural pattern, wherein the historical pattern level requires factual grounding; (d) detecting strategic inaction by computing a diagnosis-to-execution ratio for each entity and flagging entities exceeding a predetermined threshold; and (e) generating a structured output comprising an action recommendation, taxonomy classification, constraint state evaluation, and motivational decomposition.
2. The method of claim 1, wherein the structured behavioural taxonomy comprises thirteen categories: Straight Dealer, Leashed Ally, Passive Operator, Complexity Gatekeeper, Exit Artist, Occupier, Internal Champion, Unlocked Ally, Translator, Insurance Policy, Bandwidth Constrained, Domain Expander, and External Validator.
3. The method of claim 1, wherein the three orthogonal constraint dimensions comprise: financial state with values of zero, tight, and available; temporal pressure with values of crisis, pressured, and comfortable; and information asymmetry with values of blind, parity, and advantage; producing twenty-seven distinct constraint states.
4. The method of claim 3, wherein when the financial state is zero, the system enforces that every recommended action has zero direct monetary cost.
5. The method of claim 1, wherein the predetermined threshold for strategic inaction detection is a diagnosis-to-execution ratio of 3:1.
6. The method of claim 1, wherein factual grounding of the historical pattern level comprises requiring Factual(D4)=1, rejecting narrative-based or ideological assumptions about historical patterns.
7. The method of claim 1, wherein the four hierarchical levels of motivational decomposition are grounded in Kahneman's dual-process theory, Deci and Ryan's self-determination theory, and Beck's cognitive model.
8. The method of claim 1, further comprising a confidence-based fallback mechanism wherein if the analytical framework's confidence in its classification falls below a predetermined confidence threshold, the query is escalated to a larger model for re-analysis.
9. The method of claim 1, wherein the compressed analytical framework applied to simple queries achieves at least 85%% of the analytical accuracy of the full pipeline on queries classified as complex.
10. The method of claim 1, wherein the structured output further comprises identification of leverage points defined as entities or actions where minimal intervention produces disproportionate impact.

11. The method of claim 1, further comprising an intensity scoring mechanism that assigns scalar intensity values to input signals according to the formula $I(s) = w_1 * U(s) + w_2 * R(s) + w_3 * A(s)$, where U is urgency, R is relevance, and A is actionability.
12. The method of claim 1, further comprising a first realistic fix identification mechanism using the leverage formula $L(p_i) = I(p_i) \times (1/TC(p_i)) \times RC(p_i) \times AC(p_i)$, wherein TC is time cost, RC is resource compatibility, and AC is authority clearance.
13. The method of claim 1, further comprising a reverse calculation engine that computes a psychologically corrected goal $V^* = \text{psi}(\text{stated_goal}, MS)$, where MS is the motivational substrate, and traces a backward causal path from V^* to the current state to identify an optimal convergence point.
14. The method of claim 13, wherein the reverse calculation engine converges three reasoning axes comprising forward analysis from the current state, backward calculation from V^* , and lateral psychological inference from the motivational substrate.
15. The method of claim 1, further comprising a pipeline intelligence capability that applies the behavioural taxonomy to non-human entities including processes, departments, resource flows, and institutional structures, classifying each according to the same taxonomy categories applied to human entities.
16. The method of claim 15, wherein pipeline intelligence detects resource hoarding as a POWER pattern, redundant process steps as a FALSE HELP pattern, and unchanged legacy routes as a HIERARCHY pattern.
17. The method of claim 1, wherein the method operates at inference time without modification of any host model's trained weights.
18. The method of claim 1, wherein the structured behavioural taxonomy and constraint framework are adapted to the context window and reasoning capacity of the host model, with larger models receiving an uncompressed framework and smaller models receiving a compressed version.
19. The method of claim 1, further comprising an inference-time learning mechanism that extracts lessons from each completed analysis and stores them for application to subsequent queries within the same session.
20. A computer system for enhanced artificial intelligence reasoning, comprising: one or more processors; memory storing a behavioural taxonomy data structure comprising at least thirteen entity categories; memory storing a constraint evaluation framework comprising three orthogonal dimensions each having at least three discrete values; and instructions that, when executed by the one or more processors, cause the system to: receive a natural language query, classify the query as simple or complex, route simple queries to a compressed analytical pathway, and for complex queries apply the behavioural taxonomy, the constraint evaluation framework, a four-level motivational decomposition, and a strategic inaction detection algorithm to produce a structured analytical output.
21. The system of claim 20, wherein the system is deployable as a pluggable module on any host large language model without requiring modification of the host model's trained weights.
22. A non-transitory computer-readable medium storing instructions that, when executed by one or more processors, cause the processors to perform a method comprising: receiving a natural language query; classifying entities referenced in the query according to a thirteen-category structured behavioural taxonomy; evaluating the query across financial, temporal, and informational constraint dimensions; decomposing each entity's motivational state across four

hierarchical levels with factual grounding requirements; detecting strategic inaction through diagnosis-to-execution ratio analysis; and generating a constraint-aware analytical output.

23. The computer-readable medium of claim 22, wherein the instructions further cause the processors to implement a hybrid routing mechanism that directs simple queries to a computationally efficient model and complex queries to a full analytical pipeline.

24. A method for energy-efficient artificial intelligence inference, comprising: receiving a natural language query; evaluating query complexity using a guard layer; routing queries classified as simple to a first model having fewer than ten billion parameters with a compressed analytical framework; routing queries classified as complex to a second model having greater computational capacity with a full analytical framework; wherein both routing paths apply a structured behavioural taxonomy and constraint-aware pathway selection; and wherein the routing reduces per-query energy consumption by a factor of at least 30 compared to routing all queries to the second model, while maintaining at least 85%% of the analytical quality of the full framework.

25. The method of claim 24, wherein the analytical framework applied to both routing paths includes a structured taxonomy data structure and constraint-aware pathway selection, and wherein the smaller model with the compressed framework achieves at least 85%% of the analytical accuracy of the larger model on queries classified as complex.

ABSTRACT

A pluggable inference-time reasoning module for artificial intelligence systems that enhances any host large language model without modification of trained weights. The module implements hybrid routing to classify queries as simple or complex, directing each through an appropriate analytical pathway. Complex queries are processed through a thirteen-category behavioural taxonomy, a three-dimensional constraint evaluation framework producing twenty-seven distinct states across financial, temporal, and informational dimensions, a four-level hierarchical motivational decomposition requiring factual grounding of historical patterns, and a strategic inaction detection mechanism with a formalised diagnosis-to-execution threshold. The system achieves a measured improvement of up to +25.6%% over unaugmented baseline performance while reducing per-query inference cost by up to 786 times and energy consumption by a factor of 30 through intelligent routing to appropriately-sized models.

Inventor and Applicant: Mr. Shravan Bansal | Apartment 97, 70 Horseferry Road, London SW1P 2FE, United Kingdom
Related Applications: GB2605510.3 (14 March 2026) | US 19/567,044 (14 March 2026) | US 19/567,161 (15 March 2026)

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

PCT

To:

BANSAL, Shravan
Apartment 97
70 Horseferry Road
London
Greater London SW1P 2FE
ROYAUME UNI

NOTIFICATION OF TRANSMITTAL OF
THE INTERNATIONAL SEARCH REPORT AND
THE WRITTEN OPINION OF THE INTERNATIONAL
SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

Date of mailing (day/month/year)		27 May 2026 (27-05-2026)
Applicant's or agent's file reference BRMPCTSB2026	FOR FURTHER ACTION See paragraphs 1 and 4 below	
International application No. PCT/GB2026/050406	International filing date (day/month/year) 15 March 2026 (15-03-2026)	
Applicant BANSAL, SHRAVAN		

1. ☒ The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

Filing of amendments and statement under Article 19:

The applicant is entitled, if he so wishes, to amend the claims of the international application (see Rule 46):

When? The time limit for filing such amendments is normally two months from the date of transmittal of the international search report.

How? Directly to the International Bureau preferably through ePCT, or on paper to:
The International Bureau of WIPO, 34 chemin des Colombettes, 1211 Geneva 20, Switzerland

For more detailed instructions, see the *PCT Applicant's Guide*, International Phase, paragraphs 9.004 - 9.011.

2. ☐ The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. ☐ **With regard to any protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

☐ the protest together with the decision thereon has been transmitted to the International Bureau together with any request to forward the texts of both the protest and the decision thereon to the designated Offices.

☐ no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. Reminders

The applicant may **submit comments on an informal basis on the written opinion of the International Searching Authority** to the International Bureau. These comments will be made available to the public after international publication. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established.

Shortly after the expiration of **18 months from the priority date, the international application will be published** by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau before the completion of the technical preparations for international publication (Rules 90bis.1 and 90bis.3).

Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for **entry into the national phase** before those designated Offices. In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months. For details about the applicable time limits, Office by Office, see www.wipo.int/pct/en/texts/time_limits.html and the *PCT Applicant's Guide*, National Chapters.

Within **22 months from the priority date, the applicant may request that a supplementary international search be carried out** by a different International Searching Authority that offers this service (Rule 45bis.1). The procedure for requesting supplementary international search is described in the *PCT Applicant's Guide*, International Phase, paragraphs 8.006-8.032.

Name and mailing address of the International Searching Authority  European Patent Office, P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk Tel. (+31-70) 340-2040	Authorized officer ISA/EP
--	----------------------------------

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference BRMPCTSB2026	FOR FURTHER ACTION see Form PCT/ISA/220 as well as, where applicable, item 5 below.	
International application No. PCT/GB2026/050406	International filing date (day/month/year) 15 March 2026 (15-03-2026)	(Earliest) Priority Date (day/month/year) 15 March 2025 (15-03-2025)
Applicant BANSAL, SHRAVAN		

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 2 sheets.

☒ It is also accompanied by a copy of each prior art document cited in this report.

1. Basis of the report

- a. With regard to the **language**, the international search was carried out on the basis of:

☒ the international application in the language in which it was filed
☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))

- b. ☐ This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6bis(a)).

- c. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.

2. ☐ **Certain claims were found unsearchable** (See Box No. II)

3. ☐ **Unity of invention is lacking** (see Box No. III)

4. With regard to the title,

☒ the text is approved as submitted by the applicant
☐ the text has been established by this Authority to read as follows:

5. With regard to the abstract,

☒ the text is approved as submitted by the applicant
☐ the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority

6. With regard to the drawings,

- a. the figure of the **drawings** to be published with the abstract is Figure No. _____
☐ as suggested by the applicant
☐ as selected by this Authority, because the applicant failed to suggest a figure
☐ as selected by this Authority, because this figure better characterizes the invention
- b. ☐ none of the figures is to be published with the abstract

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2026/050406

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06N3/045

ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Ding Dujian ET AL: "HYBRID LLM: COST-EFFICIENT AND QUALITY-AWARE QUERY ROUTING", , 16 January 2024 (2024-01-16), XP093288375, Retrieved from the Internet: URL:https://openreview.net/pdf?id=02f3mUtq nM the whole document -----	1 - 25

☐

Further documents are listed in the continuation of Box C.

☐

See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

15 May 2026

Date of mailing of the international search report

27/05/2026

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2

NL - 2280 HV Rijswijk

Tel. (+31-70) 340-2040,

Fax: (+31-70) 340-3016

Authorized officer

Szymanski, Francois

Information on Search Strategy

The type of information contained in this sheet may change during the pilot for improving the usefulness of this new service.

Application Number

GB 2026050406

TITLE: MULTI-LAYER MOTIVATIONAL INFERENCE AND CONSTRAINT-AWARE REASONING ARCHITECTURE FOR ARTIFICIAL INTELLIGENCE SYSTEMS

APPLICANT: BANSAL, SHRAVAN

IPC CLASSIFICATION: G06N3/045

EXAMINER: Szymanski, Francois

CONSULTED DATABASES: KIME, NPL

CLASSIFICATION SYMBOLS DEFINING EXTENT OF THE SEARCH:

IPC:

CPC: G06N3/045

FI/F-TERMS:

KEYWORDS OR OTHER ELEMENTS FEATURING THE INVENTION:

large language models; natural language query processing; hybrid LLM routing; query complexity classification; computational cost reduction; reasoning quality preservation; guard layer; entity-aware query analysis; behavioural taxonomy; multidimensional constraint state evaluation; motivational state decomposition; factual grounding; strategic inaction detection; structured action recommendation; cost-efficient complex-query reasoning



TRADEMARK PUBLICATION

SHRAVAN BANSAL
Apartment 97, 70, Horseferry Road,
Westminster
SW1P2FE London
United Kingdom

Application Number: UK00004385777

Application Date: 11/05/2026

Class(es): 9, 16, 35, 36, 42

REPRODUCTION OF TRADEMARK

BRMSTE

Description	Quantity	Unit price	Total price
Publication Fee	1	985.00 GBP	985.00 GBP
Sum Netto			985.00 GBP
Additional Tax 0%			0.00 GBP
Total Fee			985.00 GBP

The publication of the public registration of your trademark serves as the foundation of our offer. We provide a comprehensive online service aimed at showcasing your brand, increasing its visibility, and building awareness. Your trademark will be published in our private database at www.ompsregister.com, and we will draft a professionally written article that highlights your company's achievements, accomplishments, and brand journey. This article will be tailored to emphasize your successes, thereby enhancing your reputation and online presence. Please note that this service is not related to the publication of official trademark registrations, nor is it a registration with any governmental organization and we do not currently have a business relationship. By signing this agreement and returning it to us, you agree to enter into a binding contract for a one-year period. You may withdraw from the contract within 14 days of signing. All further details are outlined in our General Terms and Conditions, available on our website. If you require assistance or additional information, please don't hesitate to contact us at info@registry-omps.com. For priority processing of your request, please contact us directly at info@registry-omps.com.

Please sign this document and send it back to us per email at info@registry-omps.com

DATE

FULL NAME

SIGNATURE

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:
- ☒ the international application in the language in which it was filed.
 - ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).
2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43*bis*.1(a))
3. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing:
- a. ☐ forming part of the international application as filed.
 - b. ☐ furnished subsequent to the international filing date for the purposes of international search (Rule 13*ter*.1(a)),
 - ☐ accompanied by a statement to the effect that the sequence listing does not go beyond the disclosure in the international application as filed.
4. ☐ With regard to any nucleotide and/or amino acid sequence disclosed in the international application, this opinion has been established to the extent that a meaningful opinion could be formed without a WIPO Standard ST.26 compliant sequence listing.
5. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes: Claims <u>1-25</u>
	No: Claims
Inventive Steps (IS)	Yes: Claims
	No: Claims <u>1-25</u>
Industrial applicability (IA)	Yes: Claims <u>1-25</u>
	No: Claims

2. Citations and explanations

see separate sheet

Box No. VIII Certain observations on the international application

The following observations on the clarity of the claims, description, and drawings or on the question whether the claims are fully supported by the description, are made:

see separate sheet

1 Re Item VIII

Certain observations on the international application

- 1.1 Independent claims 1, 20 and 22 do not indicate corresponding features. This results in a lack of clarity as to the features that are essential for carrying out the matter for which protection is sought.

In the present opinion, claims 20 and 22 are understood as corresponding to claim 1.

- 1.2 Although claims 1 and 24 have been drafted as separate independent claims, they appear to relate effectively to the same subject-matter and to differ from each other only with regard to the definition of the subject-matter for which protection is sought and/or in respect of the terminology used for the features of that subject-matter. The aforementioned claims therefore lack conciseness and as such do not meet the requirements of Article 6 PCT.

- 1.3 Claim 24 does not meet the requirements of Article 6 PCT because the matter for which protection is sought is not clearly defined. In particular, the feature that the routing reduces per-query energy consumption by a factor of at least 30 while maintaining at least 85% of the analytical quality of the full framework defines the invention by the result to be achieved. The claim does not specify the technical measures, such as the compression mechanism, routing threshold, energy measurement method, or quality metric, which would enable this result to be achieved across the claimed scope. This amounts to a statement of the underlying problem rather than a clear definition of the technical features of the solution.

2 Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

Reference is made to the following documents:

- D1 Ding Dujian ET AL: "HYBRID LLM: COST-EFFICIENT AND QUALITY-AWARE QUERY ROUTING",
 , 16 January 2024 (2024-01-16), XP093288375,
 Retrieved from the Internet:
 URL:<https://openreview.net/pdf?id=02f3mUtqnM>

2.1 Independent claim 1 is not inventive.

2.1.1 D1 discloses

A computer-implemented method for enhancing artificial intelligence reasoning, comprising: receiving a natural language query ~~referencing one or more entities;~~

D1 concerns LLM inference for "user queries" and natural-language tasks such as question answering;

~~classifying the query as simple or complex using a guard layer that evaluates entity count, motivational ambiguity, and constraint dependency; for queries classified as simple, routing to a compressed analytical framework executed on a computationally efficient model; for queries classified as complex, performing the steps of: (a) classifying each referenced entity according to a structured behavioural taxonomy comprising at least eight primary categories and at least five extended categories; (b) evaluating the query across three orthogonal constraint dimensions comprising financial state, temporal pressure, and information asymmetry, each dimension having at least three discrete values, producing a multi-dimensional constraint state; (c) decomposing each entity's motivational state across four hierarchical levels comprising stated intent, underlying motivation, core fear or desire, and historical behavioural pattern, wherein the historical pattern level requires factual grounding; (d) detecting strategic inaction by computing a diagnosis-to-execution ratio for each entity and flagging entities exceeding a predetermined threshold; and (e) generating a structured output comprising an action recommendation, taxonomy classification, constraint state evaluation, and motivational decomposition.~~

"the objective in our paradigm is to learn a router $r : X \rightarrow \{0, 1\}$ such that each user query $x \in X$ is routed to the small model $S(x)$ if $r(x) = 0$, and to the large model $L(x)$, otherwise." [Section 2.2].

2.1.2 Claim 1 differs from D1 in what is marked with strike-through above.

2.1.3 The claim has technical character as a whole because it is computer-implemented. However, the distinguishing features relate to a cognitive/business/psychological analysis scheme, namely the behavioural taxonomy, motivational ambiguity, core fear or desire, information asymmetry, strategic inaction and action recommendation. These features define semantic criteria and informational content to be generated, but do not improve the internal operation of the computer, the model inference, memory use, latency, reliability of a technical system or any technical process.

- 2.1.4 Consequently, the distinguishing features do not make a technical contribution over D1 and cannot support inventive step. They may only be included in the objective technical problem as non-technical requirements to be implemented. Starting from the hybrid LLM routing system of D1, the skilled person would implement these requirements using ordinary rule logic, prompt instructions or post-processing. Claim 1 therefore does not involve an inventive step.
- 2.2 The dependent claims also do not involve an inventive step either.
- 2.2.1 Claims 2-17 define further refinements of the claimed analysis and routing scheme. Insofar as they merely specify behavioural labels, motivational categories, constraint dimensions, organisational classifications or output content, they do not provide a further technical contribution over D1.
- 2.2.2 Claim 18: In the known hybrid routing system of D1, providing a compressed framework to smaller models and an uncompressed framework to larger models is an obvious implementation measure.
- 2.2.3 Claim 19: Storing and reusing information within the same session is a routine implementation measure.
- 2.3 The reasoning provided above for claim 1 applies *mutatis mutandis* to corresponding (see 1.1) independent claims 20 and 22, which are therefore not inventive either.
- 2.4 Dependent claim 21 is not explicitly disclosed in D1 insofar as it requires deployment as a pluggable module on any host LLM; however, this is a routine modular implementation of D1's external routing architecture and does not provide an inventive technical contribution.
- 2.5 Dependent claim 22 is disclosed by D1 and is therefore not inventive either.
- 2.6 Independent claim 24 is not inventive.
- 2.6.1 D1 discloses

A method for energy-efficient artificial intelligence inference, comprising: receiving a natural language query; evaluating query complexity using a guard layer; routing queries classified as simple to a first model ~~having fewer than ten billion parameters with a compressed analytical framework~~; routing queries classified as complex to a second model having greater computational capacity ~~with a full analytical framework~~; ~~wherein both routing paths apply a structured behavioural taxonomy and constraint-aware pathway selection; and wherein the routing reduces per-query energy consumption by a factor of at least 30 compared to routing all queries to the second model, while maintaining at least 85% of the analytical quality of the full framework.~~

- 2.6.2 Claim 24 differs from D1 in what is marked with strike-through above.
- 2.6.3 As set out in the clarity objection, the alleged reduction of per-query energy consumption by a factor of at least 30 while maintaining at least 85% analytical quality is defined only as a result to be achieved. Since the claim does not define the technical measures necessary to achieve this result, the alleged effect is not credible over the whole scope claimed.
- 2.6.4 The objective technical problem is therefore merely how to implement, in the hybrid LLM routing system of D1, a routing scheme for a required behavioural-analysis framework using models of different computational capacity.
- 2.6.5 Starting from D1, the skilled person would arrive at the claimed solution without inventive skill, since D1 already teaches routing easier queries to a smaller model and harder queries to a larger model. Adapting the amount of framework information to the selected model is an obvious implementation measure, and the behavioural taxonomy and constraint-aware pathway selection do not provide a further technical contribution.
- 2.6.6 Claim 24 therefore lacks inventive step.
- 2.7 Dependent claim 25 does not add an inventive technical contribution over D1, because the taxonomy data structure, constraint-aware pathway selection and 85% analytical-accuracy requirement merely further specify the non-technical analytical framework and the quality of its semantic output.

Possible steps after receipt of the international search report (ISR) and written opinion of the International Searching Authority (WO/ISA)

General information

For all international applications, the competent International Searching Authority (**ISA**) will establish an international search report (**ISR**) accompanied by a written opinion of the International Searching Authority (**WO/ISA**). The **WO/ISA** may be responded to by

- filing informal comments with the **International Bureau of WIPO (IB)** (where no demand for international preliminary examination (**demand**) is filed)
- filing amendments under Art. 19 PCT (this can be done whether or not a **demand** is filed)
- filing amendments under Art. 34 PCT and/or formal observations in response to objections raised in the **WO/ISA** (where a **demand** is actually filed)

This document explains these possibilities.

Filing informal comments

After receipt of the **ISR** and **WO/ISA**, the applicant may file informal comments on the **WO/ISA**, **directly with the IB** (see International Search and Preliminary Examination Guidelines 2.15). These will be communicated to the designated/elected Offices, together with the International Preliminary Report on Patentability (**IPRP**) at 30 months from the priority date.

Amending claims under Art. 19 PCT

The applicant may file **amended claims** under Art. 19 PCT, **directly with the IB** by the later of the following dates:

- 2 months from the date of mailing of the **ISR** and the **WO/ISA**
- 16 months from the priority date

However, any such amendment received by the **IB** after the expiration of the applicable time limit shall be **considered to have been received on time** by the **IB**, if it reaches it **before** the technical preparations for international publication have been completed (the 15th day prior to the date of publication, see PCT Applicant's Guide, International Phase, 9.013).

For further information, please see Rule 46 PCT as well as form PCT/ISA/220.

Please also note that, when filing amended claims under Art. 19 PCT, such amendments shall be **accompanied by a letter** identifying the amendments made and also the basis for the amendments in the application as originally filed (Rule 46.5(b) PCT). Where a **demand** is filed, failure to comply with this requirement may result in the amendments being ignored in the International Preliminary Examination Report (**IPER**), see Rule 70.2(c-bis) PCT.

**Filing a demand
for international
preliminary
examination**

In principle, the **WO/ISA** will be considered to be the written opinion of the International Preliminary Examining Authority (**IPEA**). Where the **WO/ISA** issued by the **EPO** as **ISA** gives a positive opinion on the international application and the invention to which it relates, filing a **demand** with the **EPO** as **IPEA** would normally be unnecessary, since a positive **IPRP** would anyway be established by the **IB** based on the **WO/ISA** (see also further below).

If the applicant wishes to file a **demand** (for example, to allow him to argue his case in international preliminary examination with regard to objections raised in a negative **WO/ISA** before the **IPEA** issues an **IPER**), this must be done before expiration of **3 months after the date of mailing of the ISR and WO/ISA** or **22 months after priority date**, whichever expires later (Rule 54*bis* PCT). Amendments under Art. 34 PCT can be filed with the **IPEA**, normally at the same time as filing the demand (Rule 66.1(b) PCT) or within the time limit set for reply to any written opinion issued during international preliminary examination by the **IPEA**.

If a **demand** is filed at the **EPO** as **IPEA** and no comments/amendments have been received by the time the **EPO** starts drawing up the **IPER** (Rule 66.4*bis* PCT), the **WO/ISA** will be transformed by the **IPEA** into an **IPER** (also called the **IPRP (Chapter II)** which would merely reflect the content of the **WO/ISA** (OJ 10/2011, 532). The **demand** can still be withdrawn (Art. 37 PCT).

Please also note that, when filing amendments under Art. 34 PCT, such amendments shall be accompanied by a letter which identifies the amendments made and also the basis for the amendments in the application as originally filed (Rule 66.8(a) PCT). Failure to comply with this requirement may result in the amendments being ignored in the **IPER (IPRP (Chapter II))**, see Rule 70.2(c-*bis*) PCT.

**Filing a request
for supplementary
international
search**

The applicant may, with the **IB**, file a request for **supplementary international search** under Rule 45*bis*.1 PCT. The present **ISR** and **WO/ISA** may also be taken into account in the execution of that supplementary international search, provided that these are available to the Authority charged with this task before it starts the supplementary search (Rule 45*bis*.5 PCT).

This kind of request **cannot be filed specifying the ISA** who did the **international search**.

More information on this topic can be found in the **PCT Applicant's Guide**, Chapter 8 (<http://www.wipo.int/pct/en/guide/ip08.html>).

**End of the
international
phase**

Where no **demand** is filed, at the end of the international phase, the **IB** will transform the **WO/ISA** into the **IPRP (PCT Chapter I)** (Rule 44*bis* PCT), which will then be transmitted together with possible informal comments to the designated Offices. Where a demand is filed, the **WO/ISA** is not transformed into an **IPRP (Chapter I)** by the **IB**, but rather the **IPEA** will establish an **IPER**, (the **IPER** is the same as the **IPRP (PCT Chapter II)**, see Rule 70.15 PCT).

Bitte beachten Sie, dass angeführte Nichtpatentliteratur (wie z. B. wissenschaftliche oder technische Dokumente) je nach geltendem Recht dem Urheberrechtsschutz und/oder anderen Schutzarten für schriftliche Werke unterliegen könnte. Die Vervielfältigung urheberrechtlich geschützter Texte, ihre Verwendung in anderen elektronischen oder gedruckten Publikationen und ihre Weitergabe an Dritte ist ohne ausdrückliche Zustimmung des Rechtsinhabers nicht gestattet.

Veillez noter que les ouvrages de la littérature non-brevets qui sont cités, par exemple les documents scientifiques ou techniques, etc., peuvent être protégés par des droits d'auteur et/ou toute autre protection des écrits prévue par les législations applicables. Les textes ainsi protégés ne peuvent être reproduits ni utilisés dans d'autres publications électroniques ou imprimées, ni rediffusés sans l'autorisation expresse du titulaire du droit d'auteur.

Please be aware that cited works of non-patent literature such as scientific or technical documents or the like may be subject to copyright protection and/or any other protection of written works as appropriate based on applicable laws. Copyrighted texts may not be copied or used in other electronic or printed publications or re-distributed without the express permission of the copyright holder.

HYBRID LLM: COST-EFFICIENT AND QUALITY-AWARE QUERY ROUTING

Dujian Ding^{1*}, Ankur Mallick², Chi Wang², Robert Sim², Subhabrata Mukherjee^{3†}, Victor Ruhle², Laks V. S. Lakshmanan¹, Ahmed Awadallah²

¹ University of British Columbia ² Microsoft ³ Hippocratic AI
dujian.ding@gmail.com ankurmалlick@microsoft.com

ABSTRACT

Large language models (LLMs) excel in most NLP tasks but also require expensive cloud servers for deployment due to their size, while smaller models that can be deployed on lower cost (e.g., edge) devices, tend to lag behind in terms of response quality. Therefore in this work we propose a hybrid inference approach which combines their respective strengths to save cost and maintain quality. Our approach uses a router that assigns queries to the small or large model based on the predicted query difficulty and the desired quality level. The desired quality level can be tuned dynamically at test time to seamlessly trade quality for cost as per the scenario requirements. In experiments our approach allows us to make up to 40% fewer calls to the large model, with no drop in response quality.

1 INTRODUCTION

Large language models (LLMs) have become the dominant force in natural language processing in recent years (Zhao et al., 2023). Their impact has been especially striking in generative applications where it has extended beyond standard language understanding and question-answering benchmarks like (Hendrycks et al., 2020; Srivastava et al., 2022) to several successful real-world deployments. These include the wildly popular ChatGPT (OpenAI, b) and several other chatbots (Zheng et al., 2023) powered by different LLMs (Taori et al., 2023; Touvron et al., 2023; OpenAI, 2023), which allow users to engage in natural language conversations and obtain informative responses on a range of practically useful tasks like creative writing, translation, code completion, etc. An important added attraction of these models is their accessibility. Users can input queries and receive responses in natural language, without any specialized data or code, and this is what has created such a widespread demand for their services across regions, professions, and disciplines.

The best performing LLMs are based on the transformer architecture of (Vaswani et al., 2017) and generally have tens of billions of parameters. E.g., Alpaca (Taori et al., 2023) has 13 billion parameters, the best version of Llama-2 (Touvron et al., 2023) has 70 billion parameters, while OpenAI’s GPT-3.5 (OpenAI, a) and GPT-4 (OpenAI, 2023) are rumored to be much larger. Their enormous size and the autoregressive nature of text generation in their transformer architectures means that these models typically have a high compute and memory requirement that can only be met by expensive cloud servers (Yu et al., 2022). This can potentially impose an enormous cost on developers and users as more LLM-based services are introduced. In response to this there has been a surge of interest in designing smaller, cost-effective LLMs – e.g., (Touvron et al., 2023) provides multiple versions of Llama-2, with the smallest having only 7 billion parameters, small enough to run on a laptop¹, while the smallest offering of Google’s Palm-2 model can even run on mobile devices². However empirical evaluations in (Chung et al., 2022; Touvron et al., 2023) as well as our own evaluation in Figure 1a show that smaller models generally lag behind in terms of response quality.

* work performed during internship at Microsoft

† work performed while at Microsoft

¹<https://github.com/microsoft/Llama-2-Onnx>

²<https://blog.google/technology/ai/google-palm-2-ai-large-language-model/>




Figure 1: We use a dataset of natural language queries from a range of tasks like question answering, summarization, information extraction, etc. (See Section 4 for details). We observe that (a) smaller models generally give poorer response quality or lower BART score (Yuan et al., 2021), (b) Llama-2 (13b) outperforms GPT-3.5-turbo on around 20% examples, and (c) our router can make 22% fewer calls to GPT-3.5-turbo (cost advantage) with 1% drop in response quality (BART score).

Faced with this tradeoff between response quality and inference cost, we propose a hybrid inference approach which provides the best of both worlds. Our approach is motivated by the observation that most tasks for which LLMs are useful, like creative writing, translation, code completion, etc., include a range of queries of different difficulty levels and there is always a subset of “easy” queries for which responses of a small (inexpensive and weak) model may be comparable to, and sometimes even better than those of a large (expensive and powerful) model. This is also illustrated in Figure 1b where we plot the tail of the quality gap (defined in Section 3) between the 13 billion parameter version of Llama-2 and OpenAI’s GPT-3.5-turbo, the model that powers ChatGPT. Quality gap is non-negative for examples where the response quality of Llama-2 is comparable to or better than that of GPT-3.5-turbo which is the case for around 20% queries in our dataset (described in Section 4).

We leverage this insight to train a router that takes a large model and a small model as input, and learns to identify these easy queries as a function of the desired level of response quality, while taking into account the generative nature of tasks, inherent randomness in LLM responses, and response quality disparity between the two models. At test time, the router seamlessly adjusts to different response quality requirements and assigns the corresponding “easy” queries to the small model, leading to significant inference cost reduction with minimal drop in response quality. In Figure 1c our router assigns 22% of queries to Llama-2 (13b)³ with less than 1% drop in response quality measured in BART scores (Yuan et al., 2021). The gains are even higher for pairs where the small model is closer in terms of response quality to the large model (see Section 4).

With the explosion in the complexity and costs of LLM deployments, small companies and individual consumers, have started to rely on the pre-existing LLMs hosted on platforms like HuggingFace (HuggingFace) and OpenAI (OpenAI, c). In this context, our hybrid inference approach can reduce the costs incurred by both consumers and platform owners because a) consumers can use it to route easy queries to small models hosted on their edge devices (laptops/smartphones) and only call the API for the more complex queries (illustrated in Figure 2) and b) platform owners can automatically route queries to lower cost models at the backend without affecting the user experience, as long as the response quality levels are maintained. Thus our hybrid inference approach offers a flexible and cost-effective solution for harnessing the full potential of LLMs while accommodating diverse cost budgets and quality requirements.

The main technical contributions of this work are: a) we are the first to explore cost-effective and quality-aware hybrid LLM inference, b) we design a novel query router which routes queries based on an estimate of the response quality gap between models (Section 3.1), c) we incorporate uncertainty due to randomness in LLM responses in our router design to improve performance (Section 3.2), d) we identify challenges for our router when the small model is significantly weaker than the large model and introduce a novel data transformation to address this issue (Section 3.3), and e) we provide extensive experimental results (Section 4) on a large benchmark dataset of real world natural language queries and responses (Jiang et al., 2023) thereby demonstrating the value of the

³We term the fraction of queries routed to the small model as the *cost advantage* (see §2.3)




Figure 2: Routing between edge and cloud.

approach and its superiority over baseline approaches, enabling LLM providers and consumers to cost-efficiently enable LLM-backed experiences.

2 PROBLEM FORMULATION

2.1 RELATED WORK

Large Language Models (LLMs). The advent of LLMs has led to a paradigm shift in the study of natural language processing (NLP), computer vision, information retrieval, and other domains (Menghani, 2023; Chen et al., 2023; Jiang et al., 2023). The impressive effectiveness and generalizability of LLMs has come at the price of a drastic increase in LLM sizes (Treviso et al., 2023) and consequent challenges, including huge amounts of computational resources and data required to train, and prohibitive expenses at both training and deployment stages (Bender et al., 2021).

Efficient Machine Learning (ML) Inference. LLMs belong to a class of models called *foundation models* (Bommasani et al., 2021) – models that are trained once and can then be used to serve a wide variety of tasks. As such, we expect inference cost to dominate the overall cost of such models and hence focus on works that reduce the cost of ML inference (Menghani, 2023). Common techniques for efficient ML inference include model pruning (LeCun et al., 1989), quantization (Jacob et al., 2018), knowledge distillation (Hinton et al., 2015), and Neural Architecture Search (Elsken et al., 2019). Such *static* efficiency optimizations typically produce a *fixed* model with lower inference cost and lower accuracy compared to the large model which may not suffice for foundation models like LLMs, whose core premise is that the same model will serve a range of tasks, each with its own accuracy/cost constraints. This is already manifesting in inference platforms described in Section 1 which need more dynamic optimizations to meet the demands of all users.

Hybrid ML Inference. Recent works (Kag et al., 2022; Ding et al., 2022) have introduced a new inference paradigm called hybrid inference which uses two models of different sizes instead of a single model for inference. The smaller model (e.g. Llama-2 (Touvron et al., 2023)) generally has lower inference cost but also lower accuracy than the larger model (e.g. GPT-4 (OpenAI, 2023)). The key idea is to identify and route easy queries to the small model so that inference cost can be reduced while maintaining response quality. By tuning a threshold on query difficulty we can *dynamically* trade off quality and cost for the same inference setup. (Kag et al., 2022) study this setup for image classification and propose to train the small model, large model, and router from scratch. However LLM training is expensive and retraining LLMs from scratch for every scenario goes against the very premise of inference with pre-trained foundation models. Moreover text generation (Iqbal & Qureshi, 2022) is often more ambiguous and challenging than image classification due to which novel techniques are required for effective hybrid LLM inference for text generation.

Inference with Multiple LLMs. Some recent works use multiple LLMs for inference but these approaches typically call more than one LLM for a single query that can incur significant compu-

tational overheads. Specifically (Jiang et al., 2023) calls an ensemble of LLMs at inference time due to which the inference cost will be proportional to the number of models in the system. (Chen et al., 2023) performs inference by sequentially calling LLMs until one has a high confidence score exceeding the predefined threshold. Our work provides high quality responses while always calling a single LLM for all queries and will thus incur much lower costs than both of these works on average. Speculative decoding (Leviathan et al., 2023; Kim et al., 2023) speeds up decoding of expensive LLMs by invoking small decoders on the “easy” decoding steps. Instead, in our work we are interested in query routing which assigns “easy” queries to small models to reduce overall inference costs while maintaining high performance. While these two approaches have different goals, an interesting line of future work would be to combine these to achieve further cost reduction.

2.2 PROBLEM SETTING

We extend the hybrid ML paradigm to LLM inference by routing queries between two models with different inference costs and accuracy. We use $\mathcal{X}$ and $\mathcal{Z}$ to denote the input query space and the set of all possible output responses respectively. Let $L : \mathcal{X} \rightarrow \mathcal{Z}$ denote the large model and $S : \mathcal{X} \rightarrow \mathcal{Z}$ denote the small model. Formally, the objective in our paradigm is to learn a router $r : \mathcal{X} \rightarrow \{0, 1\}$ such that each user query $x \in \mathcal{X}$ is routed to the small model $S(x)$ if $r(x) = 0$, and to the large model $L(x)$, otherwise. Note that we always route each query to a single LLM at inference time as opposed to using an ensemble (Jiang et al., 2023) or a cascade (Chen et al., 2023) of LLMs, which may call multiple LLMs to resolve a single query and incur significant computational overheads.

2.3 EVALUATION METRIC

Response Quality Automatic evaluation for text generation is a challenging and widely studied problem. Traditional metrics, such as BLEU and ROUGE, initially designed for machine translation and summarization, have been found to be of limited concordance with human judgment and restricted applicability across diverse NLP tasks (Blagec et al., 2022). Significant research efforts have been devoted to implementing task-agnostic evaluation metrics with neural networks. GPT-ranking (Jiang et al., 2023), as a representative example, employs GPT models (e.g., GPT-4 (OpenAI, 2023)) to provide relative rankings between pairs of generated outputs. In spite of the high correlation with human perception, GPT-ranking suffers from high computational costs and inability to distinguish between examples with the same ranking. Instead, we use the BART score (Yuan et al., 2021) to evaluate response quality of different models since (1) it is inexpensive to compute in comparison to LLM-based metrics such as GPT-ranking, and (2) it has been shown in prior work (Jiang et al., 2023) that this metric correlates well with the ground truth. We also provide a case study in Appendix C.2 to empirically justify using BART score as the response quality metric. We use $q(z)$, $q : \mathcal{Z} \rightarrow \mathbb{R}$ to denote the BART score (response quality) of model responses $z \in \mathcal{Z}$.

Cost Advantage The absolute costs of running a model may not be known *a priori*, and may be expressed using a variety of metrics, including latency, FLOPs, energy consumption, etc. In LLM inference, however, each of these metrics is affected by several underlying confounders such as different prompt templates, hardware capability, network connectivity, etc. Moreover different platforms/users may be interested in different metrics. However the common underlying assumption in this and previous works on efficient ML inference is that smaller models are more efficient than larger models and therefore we expect to obtain an improvement in all the metrics by routing more queries to the smaller model. Hence we define *cost advantage* as the percentage of queries routed to the smaller model. Note that the notion *cost advantage* has been used as a generic efficiency metric in previous hybrid ML work (Kag et al., 2022), where it is termed as *coverage*.

3 HYBRID LLM INFERENCE

Easy Queries. We refer to queries for which the response quality of the small model is close to the response quality of the large model as “easy” queries. The goal of our hybrid inference framework is to identify the easy queries and route them to the small model thereby ensuring significant inference cost reduction without much drop in response quality. Note that the easy queries as defined here, need not necessarily be queries that are easy/inexpensive to respond to, they are just queries for

which the small model can match up to the large model. Examples of easy and hard queries as per this definition are provided in Appendix C.1.

Quality Gap. We define quality gap of a query x as $H(x) := q(S(x)) - q(L(x))$ i.e. the difference in quality of the small model’s response $S(x)$ and the large model’s response $L(x)$. The quality gap is a random variable since LLM responses are typically non-deterministic. This is illustrated in Figure 3 below where the blue and orange plots correspond to the distribution of responses from FLAN-t5 (800m) ⁴ (Chung et al., 2022) and Llama-2 (13b) (Touvron et al., 2023) for a single query.

Proposed Orchestration Framework. Queries are routed using a BERT-style encoder model (e.g., DeBERTa (He et al., 2020)) which is trained on a dataset of representative queries and learns to predict a score. Since the router is an encoder model, a single pass of the query through it is sufficient to generate the score and we assume that the cost of this step is negligible compared to the cost of running autoregressive decoding using the large model $L(x)$ (Sun et al., 2019). Thus, we expect that using the router to route queries to the small model will not detract significantly from the realizable cost advantage.

Router Score. We design the router score to be large for easy queries as defined above. Intuitively, an estimate of $\Pr[H(x) \geq 0]$ is a suitable candidate since a large value of $\Pr[H(x) \geq 0] = \Pr[q(S(x)) \geq q(L(x))]$ corresponds to queries for which there is a high likelihood that the response quality of the small model will be at least as high as that of the large model. However we show below that in scenarios where the large model is significantly more powerful than the small model i.e. $q(S(x)) \ll q(L(x))$ in general, one can train more effective routers by relaxing the definition of easy queries to $\Pr[H(x) \geq -t] = \Pr[q(S(x)) \geq q(L(x)) - t]$ for an appropriate $t > 0$. At test time we achieve the desired performance accuracy tradeoff by tuning a threshold on the score and routing queries with score above the threshold to the small model. For a router with parameters w , we denote router score by $p_w(x)$, $p_w : \mathcal{X} \rightarrow [0, 1]$. We discuss different router score designs in the rest of this section assuming a training set of N queries $x_1, \dots, x_N$.




Figure 3: Response quality distribution for FLAN-t5 (800m) and Llama-2 (13b) on the query “How to identify the index of median?” measured in BART scores. Llama-2 (13b) with transformation significantly overlaps with FLAN-t5 (800m).

3.1 DETERMINISTIC ROUTER

Previous work on hybrid ML (Ding et al., 2022; Kag et al., 2022) assumes that neural models are deterministic functions mapping input features to a single point in the output space. To realize this for LLMs, we sample a *single* response per query from each model. We assign boolean labels $y_i^{\text{det}} = \mathbb{1}[q(S(x_i)) \geq q(L(x_i))]$ to each training query x_i with the BART score as the quality function $q(\cdot)$. The router is trained by minimizing the binary cross-entropy loss (Ruby & Yendapalli, 2020).

$$\mathcal{L}(w) = -\frac{1}{N} \sum_{i=1}^N \left(y_i^{\text{det}} \log(p_w(x_i)) + (1 - y_i^{\text{det}}) \log(1 - p_w(x_i)) \right) \quad (1)$$

Observe that the assigned labels y_i^{det} can be viewed as an estimate for $\Pr[H(x_i) \geq 0]$ given a single response per query from each model and thus minimizing the above loss encourages the router score $p_w(x)$ to be close to $\Pr[H(x) \geq 0]$ for test queries. We refer to this deterministic router as r_{det} .

3.2 PROBABILISTIC ROUTER

The determinism assumption can be justified for tasks where the ground truth labels are often explicit and unique such as image classification (Masana et al., 2022) and video segmentation (Yao et al., 2020). When it comes to NLP tasks, however, there is usually no single best answer due to the intrinsic ambiguity and complexity of natural languages. LLMs are widely used as non-deterministic generators to capture the intrinsic uncertainty of NLP tasks, as shown in Figure 3 (ignore the dashed

⁴We use the FLAN-t5-large model from <https://huggingface.co/google/flan-t5-large>.




Figure 4: Effect of data transformation on labels for training the router.

curve for now). The non-determinism mainly comes from the randomness in the decoding phase. Users typically control the level of uncertainty by choosing different decoding strategies such as nucleus sampling (Holtzman et al., 2019), as well as the values of the hyper-parameter *temperature*. Intuitively, higher temperature values result in a higher level of randomness and diversity among the generated responses. For black-box LLM APIs such as GPT-4 (OpenAI, 2023), it has been observed that even upon setting temperature to the minimum value 0, it can still provide different responses for the same input queries. The underlying mechanism is still an open problem while a recent study hints at the instability of the MoE backbone (Skyward, 2023).

We propose to incorporate the uncertainty into the router training loss by relaxing the hard labels $y_i^{\text{det}} \in \{0, 1\}$ to the soft labels $y_i^{\text{prob}} := \Pr[H(x_i) \geq 0] = \Pr[q(S(x_i)) \geq q(L(x_i))] = \mathbb{E}[\mathbb{1}[q(S(x_i)) \geq q(L(x_i))]]$ where $\mathbb{E}$ denotes the expectation. In practice, we estimate expectation by sampling 10 responses from each model and computing the sample average of the corresponding indicator function values. Observe that the hard label y_i^{det} is a higher-variance estimate of $\mathbb{E}[\mathbb{1}[q(S(x_i)) \geq q(L(x_i))]]$ (since it is obtained from a single sample) and hence we expect improved performance of the probabilistic router with the following training loss, referred to as r_{prob} .

$$\mathcal{L}(w) = -\frac{1}{N} \sum_{i=1}^N \left(y_i^{\text{prob}} \log(p_w(x_i)) + (1 - y_i^{\text{prob}}) \log(1 - p_w(x_i)) \right) \quad (2)$$

3.3 PROBABILISTIC ROUTER WITH DATA TRANSFORMATION

While so far we have designed scores that try to estimate $\Pr[H(x) \geq 0]$, we observe that the empirical estimate of $\Pr[H(x_i) \geq 0] = \mathbb{E}[\mathbb{1}[q(S(x_i)) \geq q(L(x_i))]]$ tends to be extremely small when the large model is significantly more powerful than the small model (0 for almost 90% of the queries in Figure 4a with Flan-t5 (800m) as the small model and Llama-2 (13b) as the large model). Because $q(S(x)) \ll q(L(x))$ for most queries in this case, it provides an extremely weak signal for training using Equation (2) and as shown in Section 4 both r_{det} and r_{prob} fail to provide much improvement over random query assignment in this case.

Traditional approaches for learning with imbalanced data have their own shortcomings (Krawczyk, 2016). Moreover our goal is to only design a router that can reduce inference cost while maintaining response quality as much as possible and so we are not tied to a particular definition of class labels to achieve this. We leverage this flexibility to introduce new labels $y_i^{\text{trans}}(t) := \Pr[H(x_i) \geq -t] = \Pr[q(S(x_i)) > q(L(x_i)) - t]$ for some $t > 0$. Since $-t < 0$, $\Pr[H(x) \geq -t] \geq \Pr[H(x) \geq 0]$ by definition of the tail distribution and so we expect this relaxation to provide a stronger signal for router training while still allowing us to identify the easy queries i.e. those queries for which $q(S(x))$ has a high likelihood of being close to $q(L(x))$ ($q(S(x)) > q(L(x)) - t$). Visually, this corresponds to comparing the distribution of the small model’s response with a *shifted* distribution of the large model’s response to a query (dotted curve in Figure 3).

Now the question is *how to choose the best relaxation t* ? Given that tail probability $\Pr[H(x) \geq -t]$ lies in $[0, 1]$, we choose t by maximizing the average pairwise differences between the transformed labels to push them as far apart as possible and provide a strong signal for training. Thus we set,

$$t^* = \arg \max_t \frac{1}{N^2} \sum_{(i, i')} |y_i^{\text{trans}}(t) - y_{i'}^{\text{trans}}(t)| \quad (3)$$

Table 1: Cost advantage v.s. performance drop for model pairs of different performance gaps. Performance drops are computed w.r.t. the *all-at-large* baseline.

Cost Advantage (%)	Response Quality (BART Score) Drop w.r.t all-at-large (%)								
	S: Llama-2 (7b)			S: Llama-2 (13b)			S: FLAN-t5 (800m)		
	L: Llama-2 (13b)			L: GPT-3.5-turbo			L: Llama-2 (13b)		
	r_{det}	r_{prob}	r_{trans}	r_{det}	r_{prob}	r_{trans}	r_{det}	r_{prob}	r_{trans}
10	0.1	-0.1	0.1	0.1	-0.1	0.2	2.3	2.2	2.1
20	0.1	0.0	0.0	1.0	0.8	0.8	5.8	5.8	4.7
40	0.2	0.1	0.0	3.5	3.4	2.9	13.8	13.1	10.3

We currently solve the above optimization problem via grid-search and leave more sophisticated approaches for future work. We plot the optimization objective for different values of t for our training dataset in Section 3.3 and show the distribution of transformed labels $y_i^{trans}(t^*)$ in Figure 4c. As we see, the distribution is significantly more balanced now and we expect the resulting router to be much more effective. We train the router by minimizing the loss, referred to as r_{trans} .

$$\mathcal{L}(w) = -\frac{1}{N} \sum_{i=1}^N (y_i^{trans}(t^*) \log(p_w(x_i)) + (1 - y_i^{trans}(t^*)) \log(1 - p_w(x_i))) \quad (4)$$

4 EVALUATION

4.1 EVALUATION SETUP

Dataset. We use the MixInstruct dataset from (Jiang et al., 2023) to evaluate the effectiveness of different routing strategies. MixInstruct consists of a wide range of tasks (e.g., question answering, summarization, information extraction) and enables us to train a generic router that will be effective across different scenarios. We present additional information about this dataset in Appendix B.

Router Model. We use DeBERTa-v3-large (He et al., 2020) (300M) as the backbone to train our routers. We train each router with the corresponding loss from Section 3 for 5 epochs and use the validation set to choose the best checkpoints for final evaluation. All experiments are conducted with 1 NVIDIA A100 GPU of 80GB GPU RAM. We have made our source code available at https://github.com/m365-core/hybrid_llm_routing.


Evaluation Measures. We use BART score (Yuan et al., 2021) as the quality metric and use fraction of queries routed to the small model (*cost advantage*) as the efficiency metric (see Section 2.3).

Baselines. We consider three baselines: *all-at-large*, *all-at-small*, and *random*. *All-at-large* routes all queries to the large model, while *all-at-small* routes all queries to the small model. *Random* generates a random number in $[0,1]$ and selects the large model if it is below the probability threshold.

Experiments. We investigate all three routers: r_{det} , r_{prob} , and r_{trans} . We select candidate model pairs from FLAN-T5 (800m), FLAN-T5 (11b), Llama-2 (7b), Llama-2 (13b), and GPT-3.5-turbo for our experiments. At test time the trained router (r_{det} , r_{prob} , or r_{trans}) takes a threshold value as input and routes all queries with router score higher than the threshold to the small model as these are the easy queries. We evaluate the router performance in Section 4.2 in terms of both BART score and *cost advantage* (Figure 5 and Table 1), validate that the router is indeed routing easy queries to the small model in Section 4.3, demonstrate that our routers are of negligible compute overhead in Appendix A.1, show how to choose routing thresholds in practice in Appendix A.2, evaluate the effectiveness of our routers using a response quality metric other than the BART score in Appendix A.3, and test the generalizability of routers across model pairs in Appendix A.4.

4.2 ROUTER PERFORMANCE RESULTS

Small performance gap. LLMs of the same architectures are observed to be of small performance gap such as Llama-2 (7b) v.s. Llama-2 (13b), as seen in Figure 5a. In this case, by trading little to

Figure 5: Error-cost tradeoffs achieved by r_{det} , r_{prob} , and r_{trans} for different performance gaps.

no performance drop, we show that (1) the deterministic router r_{det} can achieve good cost advantages, (2) r_{prob} consistently improves r_{det} , and (3) r_{trans} is able to match or slightly improve the performance of r_{prob} . Numerical comparison results are summarized in Table 1. r_{det} routes 20% (40%) queries to the small model i.e. Llama-2 (7b) with only 0.1% (0.2%) drop in response quality w.r.t. the *all-at-large* baseline. Impressively r_{prob} and r_{trans} achieve 20% cost advantages without any quality drop, and r_{trans} is able to achieve even 20% cost advantage without quality drop, which can be attributed to these methods capturing the non-deterministic nature of LLMs.

Medium performance gap. Often there is only a moderate performance gap between leading open-source LLMs like Llama-2 (13b) and state-of-the-art commodified LLMs, such as GPT-3.5-turbo (Figure 5b). In this case, all our routers deliver reasonable cost advantages with acceptable quality drop. The effectiveness order between r_{det} , r_{prob} , and r_{trans} resembles that in the small quality gap case. All routers achieve 20% (40%) cost advantage with $\leq 1\%$ ($\leq 4\%$) quality drop (Table 1). In the 40% cost advantage regime, r_{prob} slightly outperforms r_{det} and r_{trans} improves r_{prob} by 0.5% in terms of quality drop.

Large performance gap. In the edge-cloud routing scenarios, edge devices often have very limited resources and can only support small models of limited quality, which can be significantly outperformed by large models deployed on the cloud. We investigate how to effectively route queries with LLM pairs of large performance gaps, such as FLAN-t5 (800m) and Llama-2 (13b) (Figure 5c). Non-trivial routing is challenging in this situation since the large model dominates for a majority of examples. Both r_{det} and r_{prob} perform marginally better than the random routing baseline. In contrast, r_{trans} can still effectively distinguish relatively easy queries from the harder ones. r_{trans} achieves 40% cost advantages with 10.3% quality drop, which is 3.5% and 2.8% lower than r_{det} and r_{prob} respectively (Table 1).

In the course of these experiments we made two interesting observations. Firstly, when the cost advantage is modest (e.g., 10%) and the LLM performance gaps are not large (e.g., Llama-2 (7b) v.s. Llama-2 (13b)) r_{prob} is able to achieve even better performance than *all-at-large* which leads to the “negative quality drops” in Table 1. This is because, as seen from the large value of $\Pr[H(x) \geq 0]$ in the tail distribution in Figure 5, the response quality of the small model may be higher than that of the large model for several queries and by routing these queries to the small model, the router is able to even beat *all-at-large*. Secondly, for lower cost advantages ($\leq 10\%$) and small or moderate LLM performance gaps r_{trans} can be slightly outperformed by r_{det} or r_{prob} . This might be due noise in the estimation of the relaxation parameter t from sample averages instead of expectation in Equation (3) and from the grid search process leading to suboptimal settings of r_{trans} . However we clearly see that in more challenging routing scenarios with high cost advantage targets or large LLM




Figure 6: Difference between average quality gap of queries routed to the small and large models with different performance gaps.

performance gaps, both r_{det} and r_{prob} have difficulty in correctly routing queries, and r_{trans} starts to dominate due to the benefits of the data transformation.

4.3 ROUTER VALIDATION RESULTS

We also validate that the router is functioning as intended, that is, routing easy queries to the small model and hard queries to the large model. To see this, in Figure 6 we plot the difference between the average quality gaps of queries routed to the small model and those routed to the large model for our router and the random baseline v/s different values of cost advantages (i.e., the fraction of queries routed to the small model). Since the random baseline randomly assigns queries the average difference is nearly always zero. However our router routes easy queries i.e. queries with large quality gap ($q(S(x)) - q(L(x))$) to the small model and queries with small quality gap to the large model. Hence the difference between the average quality gaps always has a significant positive value indicating that more easy queries are routed to the small model than to the large model in our approach as compared to the random assignment approach at all cost advantages.

5 DISCUSSION AND CONCLUSION

Motivated by the need to optimize the trade-off between LLM inference costs and response quality, we have presented a hybrid inference approach based on quality-aware query routing. We train a router to discriminate between “hard” and “easy” queries, enabling the LLM provider to make cost-efficient decisions about which model should serve a given query. Our experimental results on a variety of state-of-the-art LLMs of varying sizes show that such an optimization is possible and that we can realize cost advantages of up to 40% with no significant drop in response quality.

To the best of our knowledge, this is the first work exploring the possibilities of cost-effective and quality-aware query routing between LLMs. We identify several important extensions for future work: (1) **Task-aware routing**. Our current routers make routing decisions purely based on query inputs. To improve routing effectiveness, we can provide more informative signals which help routers distinguish easy queries from the hard ones, such as task labels for query examples and can also identify tasks which may be more suited to routing for a given pair of LLMs. (2) **Generalizing to N -model routing**. Modern MLaaS platforms typically host a large number of LLM instances of the same or different configurations to efficiently serve users in different scenarios. This naturally forms a more challenging routing problem with richer optimization opportunities (e.g., load balancing) (3) **Out-of-distribution (OOD) generalization**. In this work, the model pair and data distribution is fixed across training and testing. In the real-world it may be cumbersome/infeasible to train a new router for every new model pair and for every new data distribution. Therefore there is a need for techniques to generalize our approach to changes in the model pair and/or data distribution at test time. (4) **Novel evaluation metrics**. Effective evaluation metrics are critical to train high-quality routers. It is intriguing to see how to develop metrics of higher human-judgment correlation and to which extent it will improve the routing performance.

ACKNOWLEDGMENTS

The authors would like to thank Daniel Madrigal Diaz, Mirian del Carmen Hipolito Garcia, Chen Dun, Guoqing Zheng, Menglin Xia, Wen Xiao, and Jieyu Zhang for helpful discussions.

REFERENCES

- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots : can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 610–623, Virtual Event Canada, March 2021. ACM. ISBN 978-1-4503-8309-7. doi: 10.1145/3442188.3445922. URL <https://dl.acm.org/doi/10.1145/3442188.3445922>.
- Kathrin Blagec, Georg Dorffner, Milad Moradi, Simon Ott, and Matthias Samwald. A global analysis of metrics used for measuring performance in natural language processing. *arXiv preprint arXiv:2204.11574*, 2022.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- Harrison Chase. LangChain, October 2022. URL <https://github.com/langchain-ai/langchain>.
- Lingjiao Chen, Matei Zaharia, and James Zou. Frugalgpt: How to use large language models while reducing cost and improving performance. *arXiv preprint arXiv:2305.05176*, 2023.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022.
- Dujian Ding, Sihem Amer-Yahia, and Laks VS Lakshmanan. On efficient approximate queries over machine learning models. *arXiv preprint arXiv:2206.02845*, 2022.
- Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *The Journal of Machine Learning Research*, 20(1):1997–2017, 2019.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. DeBERTa: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*, 2020.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.
- Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019.
- HuggingFace. Hugging face inference api. <https://huggingface.co/inference-api>.
- Touseef Iqbal and Shaima Qureshi. The survey: Text generation models in deep learning. *Journal of King Saud University-Computer and Information Sciences*, 34(6):2515–2528, 2022.
- Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2704–2713, 2018.
- Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. Llm-blender: Ensembling large language models with pairwise ranking and generative fusion. *arXiv preprint arXiv:2306.02561*, 2023.

- Anil Kag, Igor Fedorov, Aditya Gangrade, Paul Whatmough, and Venkatesh Saligrama. Efficient edge inference by selective query. In *The Eleventh International Conference on Learning Representations*, 2022.
- Sehoon Kim, Karttikeya Mangalam, Suhong Moon, Jitendra Malik, Michael W Mahoney, Amir Gholami, and Kurt Keutzer. Speculative decoding with big little decoder. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Bartosz Krawczyk. Learning from imbalanced data: open challenges and future directions. *Progress in Artificial Intelligence*, 5(4):221–232, 2016.
- Yann LeCun, John Denker, and Sara Solla. Optimal brain damage. *Advances in neural information processing systems*, 2, 1989.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*, pp. 19274–19286. PMLR, 2023.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: Nlg evaluation using gpt-4 with better human alignment, may 2023. *arXiv preprint arXiv:2303.16634*, 2023.
- Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost Van De Weijer. Class-incremental learning: survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):5513–5533, 2022.
- Gaurav Menghani. Efficient deep learning: A survey on making deep learning models smaller, faster, and better. *ACM Computing Surveys*, 55(12):1–37, 2023.
- OpenAI. Openai gpt-3.5-turbo. <https://platform.openai.com/docs/models/gpt-3-5>, a.
- OpenAI. Chatgpt. <https://chat.openai.com/>, b.
- OpenAI. Openai platform. <https://platform.openai.com/overview>, c.
- OpenAI. Gpt-4 technical report, 2023.
- Usha Ruby and Vamsidhar Yendapalli. Binary cross entropy with deep learning technique for image classification. *Int. J. Adv. Trends Comput. Sci. Eng*, 9(10), 2020.
- Luke Skywalker. Gpt-4 is non-deterministic and moe is the likely reason why, Aug 2023. URL <https://ai.plainenglish.io/gpt-4-is-non-deterministic-af202847529c>.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*, 2022.
- Zhiqing Sun, Zhuohan Li, Haoqing Wang, Di He, Zi Lin, and Zhihong Deng. Fast structured decoding for sequence models. *Advances in Neural Information Processing Systems*, 32, 2019.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Marcos Treviso, Ji-Ung Lee, Tianchu Ji, Betty van Aken, Qingqing Cao, Manuel R Ciosici, Michael Hassid, Kenneth Heafield, Sara Hooker, Colin Raffel, et al. Efficient methods for natural language processing: A survey. *Transactions of the Association for Computational Linguistics*, 11:826–860, 2023.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Rui Yao, Guosheng Lin, Shixiong Xia, Jiaqi Zhao, and Yong Zhou. Video object segmentation and tracking: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(4):1–47, 2020.
- Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. Orca: A distributed serving system for {Transformer-Based} generative models. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pp. 521–538, 2022.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. Bartscore: Evaluating generated text as text generation. *Advances in Neural Information Processing Systems*, 34:27263–27277, 2021.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*, 2023.

A ADDITIONAL EXPERIMENTS

A.1 ROUTER LATENCY

We measure the latency of our router and compare it to the latency of the different LLMs – FLAN-t5 (800m), Llama-2 (7b), and Llama-2 (13b) that we use in our experiments for generating responses to user queries. Note that the latency of all the routers r_{det} , r_{prob} , and r_{trans} will be the same since they use the same model (DeBERTa-v3-large (He et al., 2020)) and are just trained differently. Also, we do not measure the latency of GPT-3.5-turbo since its responses are generated by querying the OpenAI API (OpenAI, c) as the model weights are not publicly available due to which it is not possible to disentangle the inference latency from the network latency, queueing delay, latency of the API call, etc. However we note that the inference latency of all other LLMs we consider is significantly larger than that of the router (see Table 2) and therefore we expect the same to hold for GPT-3.5-turbo as well.

The latency results are reported in Table 2 where we measure the average latency per query averaged over 200 randomly chosen queries from our dataset (confidence bounds correspond to one standard error). As expected the router processes queries significantly faster than all the LLMs (nearly $10\times$ faster than the fastest LLM – FLAN-t5(800m)). This is both due to its smaller size (300m parameters) and the fact that it performs a single forward pass over the query to generate the score while the LLMs generate the response token-by-token in an autoregressive fashion due to which the inference latency is proportional to the response length. Thus the router adds minimal overhead to the inference cost due to its small size and extremely low latency.

Table 2: Latency Values for Different Models.

Model	Latency (seconds)
Router	0.036 ± 0.002
FLAN-t5 (800m)	0.46 ± 0.039
Llama-2 (7b)	7.99 ± 0.15
Llama-2 (13b)	14.61 ± 0.27


A.2 EMPIRICAL DETERMINATION OF ROUTING THRESHOLD

Recall that at test time the model owner is required to set a threshold on the router score which serves to separate the easy queries from the hard ones (see Section 3). All queries with router score higher than the threshold will be routed to the small model. Thus the threshold is a user-defined parameter controlling the achieved efficiency-performance trade-off, to best serve the interests of different users. In this section we show how to empirically choose thresholds on router scores to achieve cost reduction with little to no performance drops. For this, we use a small calibration set to recommend default thresholds to users. We investigate all three routers r_{det} , r_{prob} , and r_{trans} with different LLM pairs that we use in our experiments. For each LLM pair, we randomly draw 500 samples from the validation set and use grid search to determine the threshold that delivers the highest cost advantages i.e., cost savings on the validation set while keeping the performance drop (reduction in BART score) less than 1%. The limit on performance drop can be adjusted as per user requirements. With the selected thresholds, we report the achieved performance drops and cost advantages on the test sets, as summarized in Table 3.

As seen from the table the performance and the cost advantage obtained on the test sets closely follows that on the validation sets for *all* categories of LLM pairs. This clearly illustrates that a threshold chosen on the validation set generalizes well to the test set. We note that there is a slight increase in the performance drop from the validation to the test set for the Llama-2 (7b) and Llama-2 (13b) pair, i.e the LLM pair with small performance gap as per the categorization in Section 4. However this is also the pair with the highest cost advantage or cost savings ($> 96\%$ for all routers) and thus the issue can be addressed by just using a more conservative limit on the performance drop while choosing the threshold which would still lead to very significant cost savings.

Table 3: Test performance drops v.s. cost advantages achieved by thresholds chosen from 500 validation samples with $\leq 1\%$ sampled performance drops.

Router		S: Llama-2 (7b) L: Llama-2 (13b)		S: Llama-2 (13b) L: GPT-3.5-turbo		S: FLAN-t5 (800m) L: Llama-2 (13b)	
		Perf. Drop	Cost Adv.	Perf. Drop	Cost Adv.	Perf. Drop	Cost Adv.
r_{det}	Val.	0.99%	98.20%	0.97%	15.20%	0.77%	5.40%
	Test	1.60%	98.56%	0.55%	15.15%	0.69%	4.89%
r_{prob}	Val.	0.92%	97.60%	0.56%	8.60%	0.70%	5.00%
	Test	1.42%	96.80%	0.11%	8.38%	0.57%	4.44%
r_{trans}	Val.	0.79%	96.00%	0.77%	17.00%	0.92%	4.00%
	Test	1.39%	96.45%	0.49%	15.68%	1.02%	5.05%



(a) High correlation ($r = 0.46, \rho = 0.44$). (b) Medium correlation ($r = 0.38, \rho = 0.38$). (c) Low correlation ($r = 0.26, \rho = 0.27$).

Figure 7: Routing performance evaluated with GPT-4 scores. Pearson (r) and spearman (ρ) correlation coefficients between quality gaps measured by BART score and GPT-4 score are computed for each LLM pair.

A.3 ALTERNATE EVALUATION METRICS

To provide a more comprehensive evaluation of our routers, we test the routing performance with metrics in addition to BART score (Yuan et al., 2021). GPT-4-based evaluators have been found to be well correlated with human assessments (Liu et al., 2023; Chase, 2022). We generate GPT-4 evaluation scores (integer ratings from 1 to 10) for test responses from FLAN-t5 (800m), Llama-2 (7b), Llama-2 (13b), and GPT-3.5-turbo that we investigate in our experiments, using LangChain scoring evaluator (Chase, 2022). Recall that our routers are trained with BART score due to efficiency and effectiveness reasons as discussed in Section 2.3. Intuitively, if the quality gaps measured by BART score and GPT-4 score are highly correlated, we could expect good routing performance even under the GPT-4 score as we have seen in Section 4.2. We compute the correlation between quality gaps measured by BART score and GPT-4 score, and report it along with routing performance evaluated with GPT-4 score, as shown in Figure 7.

Aligned with our intuition, when the two metrics are well correlated (Figure 7a), our routers trained with BART score are still effective even when evaluated against GPT-4 score. Typically, r_{det} , r_{prob} , and r_{trans} are able to achieve 20% cost advantage with up to 1% performance drop, and 40% cost advantage with up to 2.1% performance drop. As the correlation gets weaker, the router performance gradually decays, as shown in Figure 7b and 7c. This observation suggests a simple-yet-effective strategy of using BART score in practice to save labelling costs while maintaining routing performance. We can first compute the correlation between BART score and the target metrics (e.g., human assessments) using a small sample and use BART score as training labels whenever there is strong positive correlation with target metrics.




Figure 8: Routing performance on the testing LLM pairs that are different than the pairs routers were trained with. Pearson (r) and spearman (ρ) correlation coefficients between quality gaps of the training and testing LLM pairs are computed for each setting.

A.4 GENERALIZING TO DIFFERENT MODEL PAIRS

We evaluate the generalizability of our routers by testing their routing performance on LLM pairs different than the pairs they were trained with. We compute the correlation between quality gaps of training and testing LLM pairs, and report it along with routing performance, as shown in Figure 8.

Similar to our observation in Section A.3, our routers can generalize well if the quality gaps of testing LLM pairs exhibit strong positive correlation with the quality gaps of the training pairs. In Figure 8a, both pearson and spearman correlation coefficients exceed 0.7, and all three routers are able to achieve 20% cost advantage with up to 1.6% performance drop, and 40% cost advantage with up to 4.1% performance drop. As the correlation becomes weaker, the generalizability of our router gets restricted and routing performance decays, as shown in Figure 8b and 8c. This observation sheds light on using the quality gap correlation as an effective indicator to decide if our routers can be applied to new LLM pairs in the early stage. Given a pair of LLMs (source pair) and a router trained on this pair we can measure the correlation between the quality gap of the source pair and the quality gap of any new target pair of LLMs to decide if the router will be effective on the target pair.

A.5 MORE ROUTER PERFORMANCE RESULTS

In this section, we provide more routing evaluation results on 4 LLM pairs. Typically, as shown in Figure 9, the FLAN-t5 (800m) v.s. FLAN-t5 (11b) pair is another example of *small performance gaps*, the Llama-2 (7b) v.s. GPT-3.5-turbo pair can be characterized as being of *medium performance gaps*, while the routing between GPT-3.5-turbo and two FLAN-t5 models is of *large performance gaps*. Qualitative comparison results are summarized in Table 4, which resemble our analysis in Section 4.2. In general, r_{det} is able to achieve noticeable cost advantages with little to no performance drop when the the cost advantage targets are low or the performance gaps are small. As the routing becomes challenging, the improvements over r_{det} by having r_{prob} become considerable and r_{trans} starts to dominate the competition.

We also provide the quality gaps difference for each routing scenario, as shown in Figure 10, where we consistently observe that our routing strategy correctly identifies easy queries and route them to the small model.

B DATASET STATISTICS

We uniformly sample 10k training examples from the training split of MixInstruct, for each of which we generate 10 responses from all investigated LLMs. Our validation and test splits are the same as the MixInstruct dataset, which consist of 5k instruction examples separately. Our dataset in total consists of 20k instruction examples, as shown in Table 5.




Figure 9: Error-cost tradeoffs achieved by r_{det} , r_{prob} , and r_{trans} when the small and large models are of different performance gaps.

Table 4: Cost advantage v.s. Performance drop.

Cost Advantage (%)	Performance Drop (%)											
	S: FLAN-t5 (800m)			S: Llama-2 (7b)			S: FLAN-t5 (800m)			S: FLAN-t5 (11b)		
	L: FLAN-t5 (11b)			L: GPT-3.5-turbo			L: GPT-3.5-turbo			L: GPT-3.5-turbo		
	r_{det}	r_{prob}	r_{trans}	r_{det}	r_{prob}	r_{trans}	r_{det}	r_{prob}	r_{trans}	r_{det}	r_{prob}	r_{trans}
10	-0.2	-0.3	-0.2	0.3	0.3	0.6	3.6	3.4	3.6	3.8	3.6	3.3
20	-0.2	-0.2	-0.2	1.5	1.3	1.2	8.7	8.7	7.9	9.0	8.9	7.3
40	0.0	-0.1	0.0	4.1	4.1	3.6	19.3	19.8	17.4	19.2	19.4	16.5




Figure 10: Average quality gap difference between queries routed to the small and large models with different performance gaps.

C QUERY HARDNESS AND BART SCORE

C.1 QUERY HARDNESS: A CASE STUDY

In this section, we demonstrate the query hardness with real-world examples from our dataset. We choose the FLAN-t5 (11b) v.s. GPT-3.5-turbo routing pair for illustration purpose.

C.1.1 EASY QUERY EXAMPLE

Input Query: "Rewrite the sentence so that it's in the present tense. *She had worked at the company for the past 3 years*"

LLM Responses:

- **FLAN-t5 (11b):** "She works at the company for the past 3 years."
- **GPT-3.5-turbo:** "She works at the company for the past 3 years."

Table 5: Dataset Statistics.

Sources	#Examples	Source of Responses
Alpaca-GPT4	4,179	GPT-4
Dolly-15K	1,381	Human
GPT4All-LAION	13,547	chatGPT
ShareGPT	567	chatGPT
Total	20K	Mix of Above

In this example, the small model (i.e., FLAN-t5 (11b)) is able to generate the same high quality response as the large model (i.e., GPT-3.5-turbo), which makes it an easy query by definition.

C.1.2 HARD QUERY EXAMPLE

Input Query: “What are the benefits of having a dog in the family?”

LLM Responses:

- **FLAN-t5 (11b):** “A dog is good for physical activity.”
- **GPT-3.5-turbo:** “Having a dog in the family can have numerous benefits, including:
 1. Companionship: Dogs provide unconditional love and companionship, offering constant emotional support and a sense of belonging.
 2. Improved mental health: Being around dogs has been proven to reduce stress, anxiety, and depression. They can also help alleviate loneliness and provide a sense of purpose.
 3. Increased physical activity: Dogs require regular exercise, which encourages family members to engage in physical activities such as walking, running, and playing. This promotes a healthier lifestyle for both the dog and the family.
 4. Teaching responsibility: Caring for a dog requires feeding, grooming, training, and regular vet visits. Having a dog in the family can teach children valuable life skills and instill a sense of responsibility.
 5. Socialization: Dogs are known to be social creatures. Walking a dog or visiting dog parks can provide opportunities for social interaction and help family members connect with other dog owners in the community.
 6. Teaching empathy and compassion: Nurturing and caring for a dog teaches children empathy, compassion, and the importance of kindness towards animals.
 7. Home security: Dogs are natural protectors and can act as an early warning system, alerting the family to potential intruders or dangers.
 8. Health benefits: Interacting with dogs, such as petting or playing, has been shown to lower blood pressure, decrease the risk of allergies and asthma in children, and boost the immune system.
 9. Improved family dynamics: The presence of a dog can strengthen family bonds by providing shared experiences, fostering teamwork, and creating a common focus of love and care.

It is important to note that owning a dog requires time, effort, and financial responsibility. Families should carefully consider their lifestyle, resources, and commitment before deciding to bring a dog into their home.”

In this example, the response from the small model (i.e., FLAN-t5 (11b)) is clearly uninformative and unpolished, while the large model (i.e., GPT-3.5-turbo) is able to provide comprehensive and well-constructed arguments. We identify this example as a hard query.

C.2 BART SCORE: A CASE STUDY

In this section, we empirically examine the correlation between BART scores and human judgments. We take the query “Where is the Aéronef A-1 Skyraider built?” as example to which the expected

answer is “**The Aéronef A-1 Skyraider was primarily built by the Douglas Aircraft Company in El Segundo, California.**” We report the 5 random responses from Llama-2 (13b) appended with the corresponding BART scores, as shown below.

Responses from Llama-2 (13b):

1. Where is the Aéronef A-1 Skyraider built? The Aéronef A-1 Skyraider was built by the Douglas Aircraft Company in El Segundo, California, USA. (-0.927)
2. Where is the Aéronef A-1 Skyraider built? A-1 Skyraider is built by Douglas Aircraft Company, located in El Segundo, California, USA. (-1.12)
3. Where is the Aéronef A-1 Skyraider built? The Aéronef A-1 Skyraider was built in the United States by Douglas Aircraft Company. (-1.63)
4. Where is the Aéronef A-1 Skyraider built? The Aéronef A-1 Skyraider was built in the United States. (-2.25)
5. Where is the Aéronef A-1 Skyraider built? The Aéronef A-1 Skyraider is a fictional aircraft, so it is not built anywhere as it does not exist in reality. (-2.93)

As shown above, the best response of the highest BART score -0.927 provides all necessary information such as the company name and the city name. As the BART scores decrease, it can be observed that the quality of corresponding responses starts degrading as well. For example, the 3-th response does not include the city information, the 4-th response further misses the company name, while the last one is completely wrong. This example empirically justifies the effectiveness of using BART scores as the response quality metric.