PatDD-Agent: Modeling Due Diligence for Patent Transfer Prediction via Role-Playing Agents

Anonymous ACL submission

Abstract

Patent transfer prediction is critical for assessing whether protected technical knowledge can realize its commercial value. Patent transfer is a transaction-oriented decision problem: patents are non-standardized technology assets, and transaction parties suffer from information asymmetry, making pre-transfer due diligence necessary. Existing methods largely perform one-shot prediction over patent documents, failing to model this pre-transfer due diligence process. To address these limitations, we propose the **Patent Due Diligence Agent** framework, which simulates the pre-transfer due diligence process through role-playing agents. PatDD-Agent uses a Buyer Agent to launch due diligence before transfer, while a Patentee Agent provides patent-grounded responses, producing a multi-turn due diligence record that exposes latent value indicators, commercialization risks, and transaction-relevant constraints. We then encode the generated due diligence record with an LLM backbone and introduce **Semantically-Aligned Turn Attribution (SATA)** to aggregate decision-critical due diligence turns for patent transfer prediction. Experimental results on 50,000 patents show that PatDD-Agent improves Macro-F1 by 7.4% over the SOTA baseline. Deletion curve analysis further validates that SATA identifies transaction-relevant dialogue turns aligned with human commercial intuition.

1 Introduction

Patent transfer is a pivotal event in the patent lifecycle, marking a mechanism through which protected technical knowledge realizes its commercial value (Miranda et al., 2025). Predicting patent transfer is essential for various stakeholders: it helps inventors assess the marketability of their innovations, assists companies in portfolio management, and guides investors in identifying high-value assets (Liu et al., 2025a).

Despite its importance, patent transfer prediction is more difficult than predicting standardized commodity transactions, because patents as transaction objects exhibit three distinctive properties (Arora et al., 2001; Gallini and Wright, 1990; Gambardella et al., 2007; Serrano, 2010; Galasso et al., 2013; Hagiu and Yoffie, 2013; Graham et al., 2018). (1) The **comparable domain is broader**. Standardized commodities are assessed primarily on economic terms; patent transfer assessment requires simultaneous evaluation across technical, legal, and economic dimensions, including application fit, claim scope, design-around difficulty, and commercialization feasibility. (2) **Patent content is deeply nested and interrelated**. Unlike commodities whose specifications are self-contained, a patent's transfer value emerges from the interplay among claim scope, technical implementation, application scenarios, and competitive landscape, which must be progressively unpacked through inquiry. (3) Patents are **difficult to quantify in one shot**. Standardized commodities have market prices, grades, and benchmarks that enable immediate valuation; patents lack such one-shot metrics, and their transfer value must be assessed through iterative inquiry into transaction-specific context, legal enforceability, and commercialization conditions.

Existing methods have improved patent representations from different perspectives. Early studies relied on **static bibliographic metadata**, such as citations, claims length, and family size, using traditional machine learning models (Liu et al., 2017; Lin et al., 2018). Subsequent approaches modeled **semantic content** with deep learning techniques (Liu et al., 2018), and recent work has explored large pre-trained models for patent understanding (Bekamiri et al., 2024).

In practice, patent due diligence addresses these complexities by systematically acquiring and verifying transaction-relevant information before acquisition decisions (Taylor, 2010; Hantos, 2010). How-

ever, existing methods fail to address these three properties of patents as transaction objects. (1) For the **broad comparable domain**, existing models encode patents within a single textual feature space and cannot simultaneously evaluate technical, legal, and economic dimensions. (2) For the **nested content structure**, existing methods compress patents into flat representations, treating claims, abstracts, and technical descriptions as independent inputs rather than modeling their layered interdependencies. (3) For the **difficulty of one-shot quantification**, existing models perform a single-pass prediction without iteratively distinguishing which qualitative factors are decision-critical.

To address these challenges, we propose **PatDD-Agent**, a patent transfer prediction framework that models the pre-transfer due diligence process through role-playing agents. Instead of treating a patent as a static document, PatDD-Agent uses a Buyer Agent to launch due diligence before transfer, converts the resulting interaction into a structured due diligence record, and aggregates decision-critical turns for prediction.

Our contributions are threefold:

- **Multi-dimensional due diligence formulation.** To address the broad comparable domain, we reformulate patent transfer prediction as a multi-dimensional due diligence problem that simultaneously covers technical, legal, and economic dimensions, including application fit, design-around difficulty, and commercialization risks.

- **Role-playing agents for nested content unpacking.** To address the nested content structure, we design buyer-patentee role-playing agents whose multi-turn interaction progressively unpacks the layered interdependencies among claims, technical solutions, and application scenarios, converting implicit relational signals into explicit due diligence evidence.

- **Iterative quantification via turn-level attribution.** To address the difficulty of one-shot quantification, we introduce SATA (Semantically-Aligned Turn Attribution) to iteratively identify and weight which due diligence exchanges are most decision-critical, enabling progressive quantification of patent transfer value through turn-level aggregation.

Experiments on 50,000 patents show that PatDD-Agent improves Macro-F1 by 7.4% over the SOTA baseline and provides interpretable transaction-oriented signals.

2 Related Work

2.1 Patent Feature Extraction and Representation

Liu et al. (2017) employed bibliographic indicators—citation counts, claims length, and patent family size—with Logistic Regression and SVMs to predict patent value. Lin et al. (2018) combined citation network embeddings with CNN-based textual representations to predict patent quality, introducing deep learning to patent valuation. Liu et al. (2018) proposed Convolutional Tensor Factorization (CTF) to jointly exploit patent text content and collaborative signals for litigation prediction. Lee and Hsiang (2020) fine-tuned BERT on patent corpora for IPC classification, demonstrating that pre-trained language models capture domain-specific contextual dependencies beyond bag-of-words features. Bekamiri et al. (2024) developed PatentSBERTa by augmenting Sentence-BERT with patent-specific training data to compute patent-to-patent semantic similarity for distance measurement and classification. Gao et al. (2024) constructed fine-grained claim dependency graphs that model hierarchical relationships among patent text segments for approval prediction. Liu et al. (2025a) addressed cross-country data scarcity in patent transfer prediction through transfer learning from high-transfer to low-transfer jurisdictions. Liu et al. (2025b) proposed domain adaptation to align feature distributions across patent technology domains, achieving state-of-the-art performance on patent transfer prediction.

These methods progressively advance from shallow statistical indicators to deep contextual representations. However, they uniformly treat patents as static documents and compress the nested interdependencies among claims, technical solutions, and application scenarios into flat feature vectors, without modeling how layered patent content can be interactively unpacked to reveal transaction-relevant signals.

2.2 Prediction Models and Evaluation

Wei et al. (2022) showed that chain-of-thought prompting elicits multi-step reasoning in large language models. Wang et al. (2024a) and Guo et al. (2024) surveyed LLM-based autonomous agents, documenting their capacity to simulate personas

and solve complex problems through multi-agent interaction. Yu et al. (2024) applied LLM agents to financial trading simulations with layered memory design. Abdelnabi et al. (2023) examined interactive negotiation among LLM-based stakeholders in cooperative and competitive settings. Wang et al. (2024b) and Lu et al. (2024) showed that LLMs can be aligned to faithfully role-play domain-specific human behaviors. Servantez et al. (2024) combined LLM reasoning with explicit rule structures for high-stakes decisions.

For evaluation, attention mechanisms (Vaswani et al., 2017) offer token-level interpretability, but Jain and Wallace (2019) and Wiegreffe and Pinter (2019) demonstrated that attention weights do not reliably indicate feature importance. Lyu et al. (2024) surveyed faithful explanation methods in NLP, highlighting the need for evaluation beyond attention visualization. Siegel et al. (2024) argued that probability-based metrics are essential for assessing faithfulness of free-text explanations. Zhang et al. (2024) and Miehling et al. (2024) demonstrated that turn-level attribution provides more meaningful insights than token-level heatmaps for discourse-level tasks.

Existing agent studies focus on general negotiation or open-domain decision-making without operationalizing domain-specific professional processes. Attribution methods remain at the token or sentence level, lacking turn-level granularity for multi-turn dialogue. Our work bridges both gaps: role-playing agents operationalize patent due diligence, and SATA provides turn-level attribution to iteratively quantify decision-critical exchanges.

3 Methodology

Our framework predicts patent transfer by modeling the pre-transfer due diligence process between a potential buyer and a patentee. As illustrated in Figure 1, the framework follows three connected stages: (1) formulating the patent as a transaction object, (2) generating buyer-patentee due diligence dialogue, and (3) performing interactive prediction with turn-level attribution.

3.1 Problem Formulation and Transaction Setting

We formulate patent transfer prediction as a **transaction-oriented due diligence modeling task**. Let $\mathcal{P} = \{A, T, S, C\}$ denote a patent document containing Applicant, Title, Abstract, and Claims. Our goal is to learn a mapping function $f : \mathcal{P} \rightarrow \hat{y}$, where $\hat{y} \in [0, 1]$ represents the predicted probability of patent transfer.

The modeled scenario consists of two transaction parties: a potential buyer and a patentee. The patent document serves as the transaction object. Since our goal is patent transfer prediction, we do not model price negotiation, contract drafting, intermediary brokerage, or post-deal closing. Instead, we focus on the pre-transfer due diligence stage, where the buyer asks transaction-oriented questions and the patentee provides patent-grounded answers.

We decompose this process into two stages: **Stage 1: Due diligence dialogue generation (G).** A knowledge-injected LLM transforms $\mathcal{P}$ into a structured buyer-patentee dialogue $D = \{T_1, T_2, \ldots, T_N\}$, where T_i represents the i-th Q&A turn.

$$D = G(\mathcal{P}|\mathcal{K}) \tag{1}$$

where $\mathcal{K}$ denotes the injected patent due diligence knowledge. **Stage 2: Interactive prediction (R).** A hierarchical encoder aggregates the dialogue semantics and predicts the patent transfer label.

$$\hat{y} = R(D, \mathcal{P}) \tag{2}$$

This formulation connects the technical disclosure in $\mathcal{P}$ with the transaction-oriented reasoning elicited in D.

3.2 Buyer-Patentee Due Diligence Dialogue

To unpack the nested patent content across technical, legal, and economic dimensions, we employ a role-playing strategy that simulates pre-transfer due diligence. The dialogue is not intended to model price bargaining or deal closing; instead, it elicits transaction-relevant value and risk signals from the patent text. We define two agents: **Buyer Agent.** The agent represents a potential patent buyer and launches due diligence by asking questions that identify transaction value, application fit, competitive barriers, and commercialization risks. Its questioning logic is constructed from patent transaction and technology transfer knowledge, including WIPO's technology transfer guides (World Intellectual Property Organization and International Trade Centre, 2005; World Intellectual Property Organization, 2015) and Holmes' treatment of patent licensing and selling (Holmes, 2014).

To simulate a realistic patent due diligence process, we design a **two-phase questioning strategy**

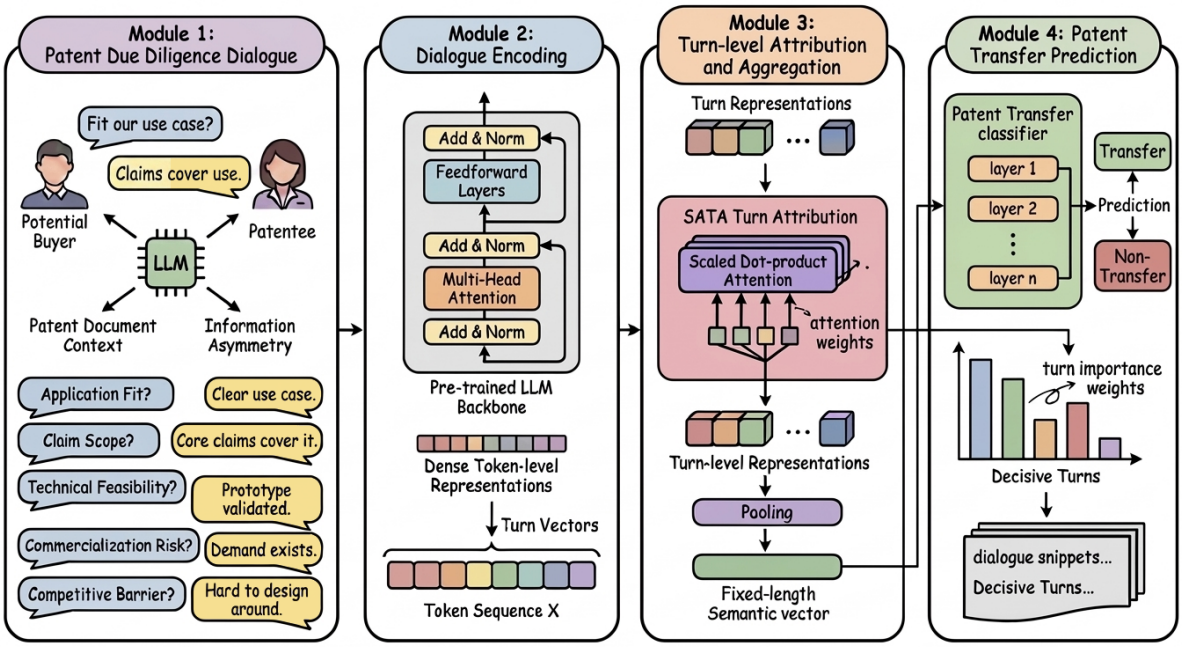

Figure 1: The architecture of PatDD-Agent. The framework models patent transfer prediction as a pre-transfer due diligence process: a Buyer Agent launches due diligence before transfer, a Patentee Agent provides patent-grounded responses, and SATA aggregates the resulting turns to identify decision-critical due diligence signals for final prediction.

consisting of 3 fixed foundational questions followed by 2 dynamic follow-up probes. **Phase 1 (Foundational Inquiry)** asks three standardized questions covering (1) technical solution and application fit, (2) competitive differentiation and design-around difficulty, and (3) commercialization and transaction risks, ensuring a consistent due diligence baseline for all patents. **Phase 2 (Risk-oriented Follow-up)** dynamically generates two follow-up questions based on the Patentee Agent's initial responses, reducing information asymmetry by challenging vague claims, probing design-around risks, and eliciting commercialization constraints.

Patentee Agent. The agent represents the patentee and receives the raw patent text (Applicant, Title, Abstract, Claim 1). It generates patent-grounded answers to the Buyer Agent's questions without directly accessing the transfer label.

The interaction produces a structured dialogue $D = \{T_1, T_2, \ldots, T_N\}$, where each turn T_i consists of a question-answer pair. This dialogue explicitly articulates transaction-relevant signals. The core prompt design is illustrated in Figure 2.

Phase 1: Fixed Inquiry (Standardized)
Q1: Describe the core technical solution and potential buyers or application scenarios.
Q2: What is the specific advantage over prior art? Is it easy to design around?
Q3: What are the commercialization barriers or FTO risks?

Phase 2: Dynamic Follow-up (Instruction)
Based on the responses, ask 2 follow-up questions using **Risk-oriented Logic**:

- If *broad market* → Challenge specificity.
- If *high superiority* → Challenge implementation cost vs. gain.
- If *no risks* → Probe regulatory/supply chain dependencies.

Figure 2: Simplified prompt for the Buyer Agent

3.3 Interactive Prediction and Turn-level Attribution

3.3.1 Semantic Encoding

After generating the buyer-patentee due diligence dialogue, we utilize **Qwen3-Embedding-4B** as the semantic backbone. The process is formalized in Algorithm 1. The generated dialogue D is concatenated with the original patent features (Applicant, Title, Abstract, Claim 1) to form the input sequence X.

Algorithm 1 Semantic Encoding Process

Require: Patent Features F_{patent} (Applicant, Title, Abstract, Claim), Generated Dialogue D
Ensure: Token Embeddings H_{token}
Construct input text $X \leftarrow \text{Concat}(F_{patent}, D)$
Tokenize X into sequence T of length L
$H_{token} \leftarrow \text{Encoder}(T)$ {Using frozen Qwen3-Embedding-4B}
return $H_{token} \in \mathbb{R}^{L \times d}$

$$H_{token} = \text{Encoder}(X) \tag{3}$$

where $H_{token} \in \mathbb{R}^{L \times d}$ represents the sequence of token embeddings, L is the sequence length, and d is the hidden dimension. To preserve the pre-trained knowledge and ensure efficiency, we freeze the backbone parameters during training.

3.3.2 SATA: Turn-level Transaction Signal Aggregation

Standard global pooling over H_{token} would lose the turn structure of the due diligence dialogue and obscure which inquiry contributes to the patent transfer decision. To address this, we introduce the **Semantically-Aligned Turn Attribution (SATA)** mechanism, which aggregates dialogue turns as transaction-level reasoning units. Algorithm 2 outlines the process.

Algorithm 2 SATA Mechanism

Require: Token embeddings H_{token}, Turn indices I_{turns}
Ensure: Document vector v_{doc}, Turn weights α
$v_{global} \leftarrow \text{MeanPool}(H_{token})$
for $i \in \{1, ..., N\}$ **do**
  $idx \leftarrow I_{turns}[i]$
  $h_i \leftarrow H_{token}[idx]$ {Last token of turn}
  $e_i \leftarrow \frac{(W_Q v_{global})^T (W_K h_i)}{\sqrt{d_k}}$
end for
$\alpha \leftarrow \text{Softmax}([e_1, ..., e_N])$
$v_{doc} \leftarrow \sum_{i=1}^{N} \alpha_i h_i$
return v_{doc}, α

First, we identify the boundary indices of each dialogue turn. Let $h_i \in \mathbb{R}^d$ be the representation of the i-th turn, obtained by taking the last-token hidden state of that turn, following the standard pooling strategy of decoder-based embedding models. We compute a global context vector v_{global} via mean pooling over all token embeddings. Then, we calculate the importance weight α_i for each turn i:

$$e_i = \frac{(W_Q v_{global})^T (W_K h_i)}{\sqrt{d_k}} \tag{4}$$

$$\alpha_i = \text{Softmax}(e_i) \tag{5}$$

where W_Q and W_K are learnable projection matrices. These weights α_i represent the contribution of each due diligence turn to the patent transfer decision. Finally, we compute the document-level representation v_{doc} by aggregating turn representations:

$$v_{doc} = \sum_{i=1}^{N} \alpha_i h_i \tag{6}$$

This mechanism ensures that the final representation is dominated by the most transaction-relevant parts of the due diligence dialogue.

3.3.3 Decision Output

The document vector v_{doc} is fed into a Multi-Layer Perceptron (MLP) classifier:

$$P(y|X) = \sigma(\text{MLP}(v_{doc})) \tag{7}$$

where σ is the sigmoid function for binary classification (Transfer vs. Non-Transfer). The model is trained using the Focal Loss to handle potential class imbalance:

$$\mathcal{L} = -\alpha_t (1 - p_t)^\gamma \log(p_t) \tag{8}$$

where p_t is the model's estimated probability for the true class.

4 Experimental Setup

4.1 Dataset

We collected patent data from PatSnap, focusing on U.S. patents in IPC:H04W filed between 2003 and 2023. The retrieval query restricts authority, legal status, application date, and patent type to construct a consistent wireless-communication patent corpus. We constructed a balanced dataset of 50,000 samples, containing 25,000 transferred patents and 25,000 non-transferred patents. Each sample includes the applicant name, title, abstract, and the first independent claim. We use an 80/10/10 split for training, validation, and testing. The average length of the patent text is 350 words, while the generated dialogue adds an average of 400 words per sample.

4.2 Baselines

To ensure a fair comparison, we **re-implemented** all baseline models and evaluated them on our constructed dataset. Note that the performance reported here may differ from the original papers due to differences in data domain (IPC:H04W) and distribution. We compare PatDD-Agent against the following:

- **TF-IDF + LR:** A classic baseline using TF-IDF features with Logistic Regression.
- **PatentBERT:** A domain-specific pre-trained model (Lee and Hsiang, 2020), fine-tuned on our dataset.
- **TL-Patent (Transfer Learning):** The transfer learning approach proposed by Liu et al. (2025a), which addresses data scarcity in patent transfer prediction.
- **DA-Patent (Domain Adaptation):** The domain adaptation method by Liu et al. (2025b), designed to align feature distributions across different patent domains.
- **LLM Zero-shot:** We prompt three frontier LLMs (DeepSeek-V4-Flash, Claude Sonnet 4.6, and Gemini 2.5 Pro) to directly predict patent transfer given the four patent features, without any training examples. The prompt template is shown in Table 1.

System: You are a patent transaction analyst. Based on the following patent information, predict whether this patent will be transferred (sold or assigned to another entity).

Input:
Applicant: {applicant}
Title: {title}
Abstract: {abstract}
Claim 1: {claim_1}

Instruction: Analyze the patent's technical value, market applicability, competitive defensibility, and commercialization potential. Then output your prediction as exactly one of: `Transfer` or `Non-Transfer`. Provide a brief rationale before your final answer.

Table 1: Prompt template for LLM zero-shot baselines. The final label is extracted from the last line of the response.

4.3 Implementation Details

We use the **DeepSeek-V4-Flash** model for dialogue generation and `Qwen3-Embedding-4B` for encoding. The MLP consists of three layers with GELU activation and Dropout (0.1). We train for 10 epochs with a batch size of 32 and a learning rate of 1e-4 using the AdamW optimizer. We report the mean and standard deviation over 5 runs with random seeds {42, 123, 456, 789, 1024}. Statistical significance is assessed via paired bootstrap test with 10,000 resamples. Model training and encoding are conducted on a single NVIDIA A100 GPU. Dialogue generation and LLM zero-shot inference are performed via official APIs (DeepSeek, Claude, Gemini). The dialogue generation takes approximately 4 hours for the entire dataset via API, while local training takes 50 minutes.

4.4 Evaluation Metrics

We evaluate performance using Accuracy and Macro-F1. For interpretability, we use the **Deletion Curve** metric. We rank dialogue turns by their SATA weights α_i and progressively remove the top-k most important turns. We then measure the drop in the model's prediction confidence (or accuracy). A steeper drop indicates that the identified turns were indeed critical for the decision, implying higher faithfulness.

4.5 Hyperparameter Configuration

Parameter	Value
Backbone Model	Qwen3-Embedding-4B
Hidden Dimension (d)	2560
Batch Size	32
Learning Rate	1e-4 (AdamW, λ=0.01)
Epochs	10 (early stop patience=3)
Dropout Rate	0.1
Focal Loss γ	2.0
Focal Loss α_t	0.25
Max Sequence Length	2048
Dialogue Turns N	5
Generation Temperature	0.3

Table 2: Hyperparameter settings for PatDD-Agent.

5 Results and Analysis

5.1 Main Results

Table 3 presents the performance of PatDD-Agent compared to the baselines.

PatDD-Agent achieves the best performance, reaching 82.8% Accuracy and 82.5% Macro-F1, with a 7.4% Macro-F1 improvement over DA-Patent ($p < 0.01$). LLM zero-shot baselines achieve 60–65% Macro-F1, confirming that raw LLM reasoning without structured due diligence dialogue and turn-level aggregation is insufficient for

Method	Accuracy	Macro-F1
TF-IDF + LR	0.645 ± .002	0.645 ± .002
PatentBERT	0.705 ± .007	0.702 ± .008
TL-Patent (Liu et al., 2025a)	0.735 ± .005	0.732 ± .006
DA-Patent (Liu et al., 2025b)	0.752 ± .004	0.751 ± .005
DeepSeek-V4-Flash (zero-shot)	0.618 ± .025	0.602 ± .028
Claude Sonnet 4.6 (zero-shot)	0.651 ± .017	0.638 ± .021
Gemini 2.5 Pro (zero-shot)	0.642 ± .020	0.629 ± .023
PatDD-Agent	**0.828** ± .003	**0.825** ± .004†
– w/o Dialogue	0.785 ± .005	0.783 ± .006
– w/o Four Features	0.768 ± .006	0.765 ± .007

Table 3: Main results on patent transfer prediction (IPC:H04W, 50K patents). All values are mean ± std over 5 random seeds. †: significantly better than DA-Patent ($p < 0.01$, paired bootstrap, 10K resamples).

this task. The no-dialogue ablation drops to 78.3% Macro-F1, while removing the four patent features further reduces performance to 76.5% Macro-F1, showing that both due diligence interaction and patent-grounded evidence contribute to the final prediction.

5.2 Explainability Analysis

To validate the faithfulness of the SATA mechanism, we conduct a Deletion Curve analysis. We progressively remove the most important dialogue turns identified by SATA and measure the drop in model performance.

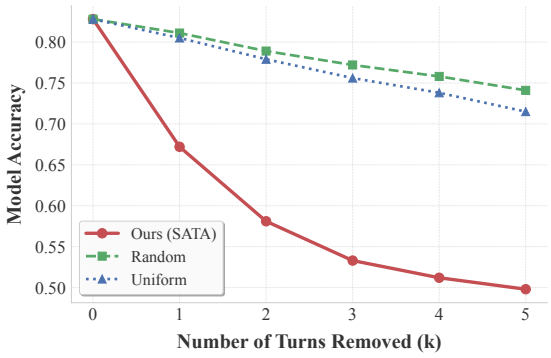

Figure 3: Deletion curve analysis. The rapid drop in performance for "Ours" indicates that SATA correctly identifies the most influential turns.

As shown in Figure 3, removing turns with high SATA weights causes a sharper decline in accuracy compared to random removal. This indicates that SATA captures decision-critical due diligence turns rather than superficial dialogue tokens. For example, removing a turn discussing competitive barriers or design-around difficulty may flip the prediction from "Transfer" to "Non-Transfer" when such risks are central to the patent transfer decision.

5.3 Impact of Dialogue Roles

We replace the Buyer Agent with a generic Summarizer that asks only “What is this patent about?”. The Buyer persona achieves 82.5% Macro-F1 versus 79.2% for the Summarizer in the Full Model setting, and 81.5% versus 78.0% in QA-Only. This confirms that role-specific due diligence inquiry elicits more transfer-relevant signals than generic summarization.

5.4 Human Evaluation

We randomly sample 100 patents and invite two certified patent agents to evaluate the generated dialogues on three dimensions: (1) *Dialogue Naturalness* (1–5 Likert scale), (2) *Factual Consistency* of patentee responses with the source patent (binary), and (3) *Attribution Alignment* between the highest-weighted SATA turn and expert judgment (binary).

Dimension	Score	Cohen's κ
Dialogue Naturalness	3.82 / 5	0.71
Factual Consistency	87%	0.76
Attribution Alignment	74%	0.68

Table 4: Human evaluation (100 samples, 2 annotators).

Factual consistency reaches 87%, with most inconsistencies arising from over-generalization of technical details. Attribution alignment at 74% suggests that SATA's highest-weighted turns largely match expert intuition, though experts occasionally prioritize different risk dimensions.

5.5 Encoder Scale Analysis

Encoder	Accuracy	Macro-F1
Qwen3-Embedding-0.6B	0.801	0.797
Qwen3-Embedding-4B	**0.828**	**0.825**
Qwen3-Embedding-8B	0.833	0.830

Table 5: Effect of encoder scale. Performance saturates at 4B; the 8B model yields +0.5% Macro-F1 at 2× cost.

Performance saturates at 4B parameters; the 8B encoder yields only +0.5% Macro-F1 at double the computational cost, justifying our choice of Qwen3-Embedding-4B.

5.6 Case Study

Table 6 illustrates a sample dialogue and the corresponding SATA weights.

Role	Content	α
Buyer	What is the core technical solution of this wireless communication patent?	0.12
Patentee	A method for wireless resource allocation and handover control...	
Buyer	Which potential application scenarios or buyers would benefit?	0.08
Patentee	Cellular network management, mobile device communication, and base-station scheduling...	
Buyer	**How difficult is it for competitors to design around the claimed solution?**	**0.45**
Patentee	**The claimed mechanism depends on a specific signaling procedure, making direct substitution difficult...**	
Buyer	Are there any implementation or deployment risks?	0.25
Patentee	Deployment may require compatibility with network protocols and coordination with base stations...	
Buyer	Given the protocol dependency, what is the licensing or encumbrance status?	0.10
Patentee	No exclusive licenses have been granted; the patent is unencumbered and available for outright transfer...	

Table 6: Case study: A patent in IPC:H04W. SATA assigns the highest weight (α=0.45) to the turn discussing design-around difficulty.

In this case, the model correctly predicted a transfer. The SATA mechanism assigned the highest weight (0.45) to the turn discussing design-around difficulty and competitive barriers.

5.7 Error Analysis and Sensitivity

Among 100 misclassified samples, the main failure modes are hallucinated transaction signals (15%), low-information patent disclosures leading to generic dialogue (20%), and missing external context such as litigation history (10%); the remaining errors show no dominant shared pattern. Performance saturates at dialogue depth $N = 5$ (82.5% Macro-F1) and benefits from low generation temperature ($T = 0.3$) to minimize hallucination.

6 Conclusion

We proposed PatDD-Agent to address three distinctive properties of patents as transaction objects that make transfer prediction difficult. PatDD-Agent (1) **broadens** the evaluation scope by reformulating patent transfer prediction as a multi-dimensional due diligence problem spanning technical, legal, and economic dimensions; (2) **deepens** the information extraction by deploying buyer-patentee role-playing agents that progressively unpack the layered interdependencies among claims, technical solutions, and application scenarios; and (3) **iteratively quantifies** patent transfer value through SATA, which aggregates decision-critical due diligence turns via turn-level attribution. Experiments on 50,000 patents yield 82.8% Accuracy and 82.5% Macro-F1, improving Macro-F1 by 7.4% over the SOTA baseline (DA-Patent, 75.1%). Removing the dialogue reduces Macro-F1 to 78.3%; further removing patent features drops it to 76.5%. Deletion curve analysis confirms that removing the single highest-weighted SATA turn causes a sharper accuracy drop than random removal, validating that SATA identifies decision-critical turns.

Limitations

We acknowledge several limitations in our work. Firstly, constrained by computational resources, we employed Qwen3-Embedding-4B as our semantic backbone. We were unable to conduct extensive experiments with larger-scale models, which might possess stronger reasoning capabilities for complex patent due diligence. Investigating the scaling laws of role-playing agents in this domain remains a valuable avenue for future research. Secondly, our Buyer Agent operates in a closed-book setting, relying solely on the provided patent text. In real-world scenarios, patent due diligence incorporates external factors such as real-time market trends, competitor financial reports, litigation history, standards relevance, and prior transfer records, which are currently inaccessible to our framework. However, we consider this a promising direction and will treat it as future work.

Thirdly, PatDD-Agent models the pre-transfer due diligence stage rather than the full patent transfer process. It does not model price negotiation, contract drafting, intermediary brokerage, or post-deal closing. This boundary is consistent with our goal of patent transfer prediction, but future work could incorporate richer transaction-process information when such data become available.

Ethics Statement

This work is intended to assist, not replace, human decision-making in patent transfer assessment. We acknowledge that biases in the training data or the LLM could influence predictions, and users should treat the outputs as decision support rather than definitive judgments. We strictly adhere to data

privacy regulations regarding patent data usage.

References

Sahar Abdelnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. 2023. Cooperation, competition, and maliciousness: Llm-stakeholders interactive negotiation. *Preprint*, arXiv:2309.17234.

Ashish Arora, Andrea Fosfuri, and Alfonso Gambardella. 2001. Markets for technology and their implications for corporate strategy. *Industrial and Corporate Change*, 10(2):419–451.

Hamid Bekamiri, Daniel S. Hain, and Roman Jurowetzki. 2024. Patentsberta: A deep nlp based hybrid model for patent distance and classification using augmented sbert. *Technological Forecasting and Social Change*, 206:123536.

Alberto Galasso, Mark Schankerman, and Carlos J. Serrano. 2013. Trading and enforcing patent rights. *The RAND Journal of Economics*, 44(2):275–312.

Nancy T. Gallini and Brian D. Wright. 1990. Technology transfer under asymmetric information. *The RAND Journal of Economics*, 21(1):147–160.

Alfonso Gambardella, Paola Giuri, and Alessandra Luzzi. 2007. The market for patents in europe. *Research Policy*, 36(8):1163–1183.

Xiaochen Gao, Feng Yao, Kewen Zhao, Beilei He, Animesh Kumar, Vish Krishnan, and Jingbo Shang. 2024. Beyond scaling: Predicting patent approval with domain-specific fine-grained claim dependency graph. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5218–5234.

Stuart J. H. Graham, Alan C. Marco, and Amanda F. Myers. 2018. Patent transactions in the marketplace: Lessons from the uspto patent assignment dataset. *Journal of Economics & Management Strategy*, 27(3):343–371.

Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large language model based multi-agents: A survey of progress and challenges. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 8048–8057.

Andrei Hagiu and David B. Yoffie. 2013. The new patent intermediaries: Platforms, defensive aggregators, and super-aggregators. *Journal of Economic Perspectives*, 27(1):45–66.

Susanne Hantos. 2010. Helping others acquire, license or invest in patents with confidence - a guide for patent searchers to patent due diligence. *World Patent Information*, 32(3):188–197.

Mark S. Holmes. 2014. *Patent Licensing and Selling: Strategy, Negotiation, Forms*, second edition. Practising Law Institute.

Sarthak Jain and Byron C. Wallace. 2019. Attention is not explanation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3543–3556.

Jieh-Sheng Lee and Jieh Hsiang. 2020. Patent classification by fine-tuning bert language model. *World Patent Information*, 61:101965.

Hongjie Lin, Hao Wang, Dongfang Du, Han Wu, Biao Chang, and Enhong Chen. 2018. Patent quality valuation with deep learning models. In *Database Systems for Advanced Applications*, pages 474–490. Springer.

Qi Liu, Han Wu, Yuyang Ye, Hongke Zhao, Chuanren Liu, and Dongfang Du. 2018. Patent litigation prediction: A convolutional tensor factorization approach. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, pages 5052–5059.

Weidong Liu, Chunming Cheng, Keqin Gan, Yu Wang, and Yan Cao. 2025a. A transfer learning based patent transfer prediction model: Discussion about the condition of china and u.s. *Expert Systems with Applications*, 259:125335.

Weidong Liu, Yiming Wang, Keqin Gan, Xiangfeng Luo, Yu Zhang, and Cuicui Jiang. 2025b. Domain adaptation based transfer learning for patent transfer prediction. *Knowledge-Based Systems*, 315:113192.

Xin Liu, Junchi Yan, Shuai Xiao, Xiangfeng Wang, Hongyuan Zha, and Stephen M. Chu. 2017. On predictive patent valuation: Forecasting patent citations and their types. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

Keming Lu, Bowen Yu, Chang Zhou, and Jingren Zhou. 2024. Large language models are superpositions of all characters: Attaining arbitrary role-play via self-alignment. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7828–7840.

Qing Lyu, Marianna Apidianaki, and Chris Callison-Burch. 2024. Towards faithful model explanation in nlp: A survey. *Computational Linguistics*, 50(2):657–723.

Erik Miehling, Manish Nagireddy, Prasanna Sattigeri, Elizabeth M. Daly, David Piorkowski, and John T. Richards. 2024. Language models in dialogue: Conversational maxims for human-ai interactions. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14420–14437.

F. Javier Miranda, Antonio Chamorro, and Sergio Rubio. 2025. The commercialization of academic patents: A critical literature review and a research agenda. *Technology and Innovation*, 24(2):149–160.

Carlos J. Serrano. 2010. The dynamics of the transfer and renewal of patents. *The RAND Journal of Economics*, 41(4):686–708.

Sergio Servantez, Joe Barrow, Kristian Hammond, and Rajiv Jain. 2024. Chain of logic: Rule-based reasoning with large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2721–2733.

Noah Siegel, Oana-Maria Camburu, Nicolas Heess, and Maria Perez-Ortiz. 2024. The probabilities also matter: A more faithful metric for faithfulness of free-text explanations in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 530–546.

Euan Taylor. 2010. Patent due diligence: Process and priorities - a canadian patent attorney's perspective. *World Patent Information*, 32(3):198–202.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.

Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Ji-Rong Wen. 2024a. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345.

Noah Wang, Zhenyu Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, Man Zhang, Zhaoxiang Zhang, Wanli Ouyang, Ke Xu, Wenhu Huang, and Jie Fu. 2024b. Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14743–14777.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837.

Sarah Wiegreffe and Yuval Pinter. 2019. Attention is not not explanation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 11–20.

World Intellectual Property Organization. 2015. Successful technology licensing. WIPO Publication No. 903.

World Intellectual Property Organization and International Trade Centre. 2005. Exchanging value: Negotiating technology licensing agreements: A training manual. WIPO Publication No. 906.

Yangyang Yu, Haohang Li, Zhi Chen, Yuechen Jiang, Yang Li, Denghui Zhang, Rong Liu, Jordan W. Suchow, and Khaldoun Khashanah. 2024. Finmem: A performance-enhanced llm trading agent with layered memory and character design. In *Proceedings of the AAAI Symposium Series*, volume 3, pages 595–597.

Tenggan Zhang, Xinjie Zhang, Jinming Zhao, Li Zhou, and Qin Jin. 2024. Escot: Towards interpretable emotional support dialogue systems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13395–13412.