

IFC-Bank of Italy Workshop on "Data science in central banking"

18-20 February 2025

AI-driven legal summarization for financial regulation monitoring¹

Diana López Avilés and Isidora Díaz Jorquera,
Central Bank of Chile

¹ This contribution was prepared for the workshop. The views expressed in this publication are those of the authors and do not necessarily represent the official views of the Committee, its members, or the BIS.

AI-Driven Legal Summarization for Financial Regulation Monitoring

Diana López Avilés¹, Isidora Díaz Jorquera²

Abstract

In Chile, financial regulation is becoming increasingly intricate due to the rapid expansion of the digital payments ecosystem. This growth has brought a surge of new market participants, evolving business models, and changing competitive dynamics — all of which contribute to more frequent and complex legal disputes. As a result, regulatory bodies, including the Central Bank of Chile, must navigate a growing volume of nuanced legal material to monitor compliance and actively oversee emerging market trends. To address this challenge, this article proposes a generative AI application that automates the summarization of legal arguments in cases involving payment and banking disputes, using public daily status data from Chile’s Court of Defense of Free Competition. By fine-tuning mBART-50, a multilingual transformer model, the tool produces specialized one-page summaries that reduce the Central Bank’s manual workload, supporting the early identification of market pressure points and providing insights into potential shifts in the regulatory environment — thereby strengthening financial oversight and enabling more agile policy responses.

Keywords: Legal summarization, GenAI, Financial regulation, Digital payments

JEL classification: C45, G28, K22

Content

1. Introduction	2
2. Data	4
3. Methodology.....	6
Data Preprocessing.....	6
Model Selection	6
Fine-Tuning Strategy.....	7
Evaluation	8

¹ Central Bank of Chile (dlopez@bcentral.cl)

² Central Bank of Chile (idiiaz@bcentral.cl)

The views are those of the authors and do not necessarily reflect the views of the Central Bank of Chile or its board members. We thank Rosario Celedón, Gabriel Aparici and Pablo Furche for their valuable comments and suggestions.


4. Results.....	8
5. Concluding remarks.....	11
6. References	12

1. Introduction

The landscape of digital payments in Chile has undergone significant transformation in the past decade, with rapid growth in both number and amount of transactions. According to the Central Bank of Chile's *Informe de Sistemas de Pago 2024* [1], digital payment activity—including both card-based payments and electronic fund transfers—has grown steadily, reflecting a nationwide move away from cash-based transactions. As of early 2024, the annual volume of digital payments surpassed six billion transactions, marking a substantial increase from the previous year (Graph 1). This equates to an average of over 300 digital payments per person, underscoring the maturity of Chile's transition to a more integrated and accessible digital financial infrastructure [1].

Digital payments in Chile

Graph 1



* Credit and debit card data include both businesses and individuals. TEF (Electronic Funds Transfers) data only account for individuals and include both intrabank and interbank transactions.

Sources: Central Bank of Chile based on data from the Financial Market Commission (CMF) [1]

Multiple structural factors have contributed to this shift. The broadening availability of payment instruments and the continued rollout of point-of-sale (POS) systems by both traditional and emerging actors have significantly improved the user experience. Meanwhile, the ecosystem of electronic transfers has enabled seamless peer-to-peer (P2P) payments, with these transfers accounting for a disproportionately large share of transaction value despite representing a smaller portion of total transactions [1]. By contrast, person-to-business (P2B) transfers via EFT channels remain less prevalent, pointing to ongoing gaps in merchant integration.

Payment card networks have generally led the way in terms of interoperability and merchant reach. These schemes benefit from legacy arrangements among issuers, acquirers, and network providers, which allow for coordinated rule-setting and the early adoption of new technologies such as contactless payments, QR code integration, and mobile wallets. This structural cohesion contrasts with the more fragmented development of EFT systems in the P2B segment, where industry-wide standards and real-time merchant interfaces are still evolving.

However, the rapid expansion of the digital payments sector has introduced a new layer of institutional and legal complexity. The growing involvement of fintechs, cross-border acquirers, and payment service providers (PSPs) has led to disputes over fee structures, network participation rules, and operational limitations. These tensions have increasingly been brought before Chile's Tribunal for the Defense of Free Competition (TDLC) and other oversight bodies, as participants seek clarity on competitive conduct and market access. Recent decisions by the Supreme Court have addressed these concerns by establishing clearer guidelines on cross-border acquiring, pricing transparency, and non-discriminatory card acceptance. In doing so, the Court has effectively shifted some aspects of private network governance into the public regulatory domain [2].

Regulators have responded accordingly. The Central Bank of Chile recently expanded its regulatory framework to oversee PSPs and cross-border acquiring arrangements, including bringing certain actors under the early supervision of the Financial Market Commission (CMF) to mitigate operational risks to merchants. Simultaneously, the CMF has launched implementation of the Open Finance System (SFA), which will —over a multi-year horizon— integrate third-party providers such as payment initiators into the financial ecosystem, enabling them to initiate transfers on behalf of users with proper authorization [3]. These reforms represent a step toward greater modularity in financial services, as well as more active oversight of previously under-regulated business models.

In parallel, these regulatory shifts have increased the informational demands placed on oversight institutions. The BCCh, in particular, is tasked under its Organic Constitutional Law with safeguarding the integrity of both domestic and international payments. Fulfilling this mandate requires real-time awareness of legal and institutional developments in the market. Yet the volume and complexity of legal proceedings —particularly those involving competition law or the introduction of novel financial technologies— pose challenges for timely analysis. Documents submitted to the TDLC and similar bodies are frequently lengthy, domain-specific, and difficult for non-specialist audiences to fully grasp.

Moreover, given the underdevelopment of certain market sectors, it is likely that frictions between incumbent actors and new entrants will re-emerge as activity expands into these areas. While some disputes have been addressed through recent regulatory and judicial decisions, evolving frameworks and diversified business models may reactivate tensions around access, pricing, and compliance. This prospect reinforces the critical role of systematic market monitoring in financial regulation.

To address this challenge, the remainder of the article presents a natural language processing tool designed to automate the summarization of legal arguments in disputes related to means of payment. The tool, based on a fine-tuned multilingual translation model (mBART-50), ingests raw TDLC case documents, and

generates concise, one-page summaries that highlight key claims, involved parties, and procedural context, with the primary objective to streamline regulatory review processes, reduce manual workload, and strengthen the Central Bank’s capacity to proactively identify market stress points and emerging regulatory needs.

2. Data

The dataset encompasses 19 cases related to payments brought before the TDLC, covering the period from 2016 to 2024. Each case includes a minimum of two documents (typically a claim and counterclaim) and an average of four, each ranging from 40 to 60 pages in length. The dataset covers both contentious cases—involving adversarial disputes and active judicial intervention—and non-contentious cases, which consist of petitions submitted by individual entities where judicial review occurs only as mandated by law.

The complexity of each document varies considerably, shaped primarily by the number and type of entities involved, as well as the depth and intricacy of their claims, responses, and counterclaims. These documents often contain dense legal arguments, detailed business model descriptions, economic justifications, and references to regulatory frameworks, with considerable variation in scope and emphasis depending on the specifics of each case (Graph 2.1).

Original document example (extract)

Graph 2.1

H. TRIBUNAL DE DEFENSA DE LA LIBRE COMPETENCIA Luis Eduardo Toro Bossay, Francisco Bórquez Electorat, José Luis Corvalán Pérez y Catalina Villalobos Hinojosa, abogados, en representación de Servicios Visa International Limitada, en estos autos caratulados “Demanda de Ebanx Chile Ltda. en contra de Servicios Visa Internacional Ltda.”, Rol C N° 444-2022, al H. Tribunal de Defensa de la Libre Competencia (“H. Tribunal” o “H. TDLC”) respetuosamente decimos: De conformidad con lo dispuesto en los artículos 20 y 29 del Decreto Ley N° 211 (“DL 211”) y el artículo 308 del Código de Procedimiento Civil, contestamos la demanda interpuesta por Ebanx Chile Limitada (“Ebanx” o “Demandante”), en contra de nuestra representada, solicitando desde ya al H. Tribunal su rechazo en todas sus partes, con costas, en virtud de los antecedentes que pasamos a exponer. En esta presentación, se verá con claridad la serie de imprecisiones y contradicciones de la demanda que Ebanx ha interpuesto en contra de nuestra representada, en donde ha intentado argumentar que la Regla de Restricción en Adquirencia Transfronteriza de Visa (“Regla de Adquirencia Transfronteriza” o la “Regla”) de Visa constituye un abuso de posición dominante, cuando en la práctica, ello no puede distar más de la realidad. En efecto, la Demandante ha acompañado información incompleta para efectos de crear	
--	--

Source: TDLC

The dataset, while substantial, presents some inherent limitations. First, since the analysis focuses solely on cases publicly available through the TDLC platform, it may not capture the full spectrum of disputes within the payments ecosystem. For

Estado Diario

Causa: C-444-2022 D Demanda de Ebanx Chile Ltda. en contra de Servicios Visa Internacional Ltda.
Causa(s) Acumulada(s): C-450-2022 D; C-448-2022 D;

Cuaderno: Cuaderno principal Fecha Ingreso: 06/05/2022

Expediente

Información de Causa

Partes

Trámite	Descargar	Detalles	Referencia	Fecha	Parte	Firmantes	Foja
Resolución			Ha lugar	17-12-2024			1151 - 1151
Escrito			Reposición	02-12-2024	DLOCAL CHILE SpA		1150 - 1150
Resolución			Ha lugar	26-11-2024			1149 - 1149
Escrito			Solicita lo que indica	07-06-2024			1113 - 1113
Escrito			Solicita lo que indica	07-06-2024	DLOCAL CHILE SpA		1114 - 1114
Escrito			Solicitud que indica	07-06-2024	Ebanx Chile Limitada		1118 - 1118
Escrito			Solicitud que indica	07-06-2024	Ebanx Chile Limitada		1116 - 1116
Escrito			Solicitud que indica	07-06-2024	Servicios Visa Internacional Limitada		1122 - 1122
Escrito			Solicitud que indica	07-06-2024	PAYU CHILE S.A.		1120 - 1120
Actuación			Certificación_notifica resolución que indica	15-11-2024			1148 - 1148
Resolución			Se fija nueva fecha de audiencia de conciliación	14-11-2024			1147 - 1147
Escrito			Solicita lo que indica	12-11-2024			1146 - 1146
Actuación			Notifica resolución que indica	05-09-2024			1145 - 1145
Resolución			Suspensión del procedimiento de común acuerdo	03-09-2024			1144 - 1144
Escrito			Solicita suspensión de procedimiento	30-08-2024			1142 - 1142
Resolución			Téngase presente	02-09-2024			1143 - 1143
Escrito			Delega poder	28-08-2024	DLOCAL CHILE SpA		1140 - 1140

Source: TDLC


example, over the past decade, most cases have involved interchange fees and acquirers' costs (Graph 2.2).

Second, the diversity in document formats—ranging from scanned PDFs to structured digital filings—introduces variability in data quality and processing reliability (Graph 2.3). To mitigate the risk of overfitting on such a small and heterogeneous dataset, we prioritized model selection over extensive fine-tuning, relying on a fully pre-trained multilingual model capable of handling varied input formats and generalizing across cases.

The dataset remains dynamic and continues to evolve, with new documents frequently uploaded to the TDLC platform as long as cases remain open. These ongoing updates reflect the need for continuous monitoring to remain aligned with real-time judicial developments.

Top TF-IDF Words Across Paragraphs

Graph 2.2



Source: Authors' own elaboration based on selected cases from the TDLC

3. Methodology

Data Preprocessing

Each PDF document undergoes a structured preprocessing pipeline aimed at creating high-quality, informative units suitable for supervised fine-tuning. Initially, the text is extracted and segmented into individual paragraphs based on structural cues such as spacing, indentation, and formatting patterns. Subsequently, shorter paragraphs are merged into coherent textual chunks with a minimum length of 600 characters to ensure sufficient content for a stand-alone summary³. Overall, the final dataset contains 3,524 chunk-summary pairs, with an average of 50 words each.


This segmentation ensures that each resulting chunk is contextually meaningful and sufficiently informative, optimizing the process for manual summarization. By creating adequately sized textual units, annotators could more reliably produce high-quality summaries, improving the overall relevance of each chunk-summary pair used for training. The granularity achieved through this preprocessing facilitates detailed scrutiny, enabling downstream models to effectively focus only on critical arguments, content and identify nuanced differences between perspectives.

Model Selection

A multilingual transformer architecture, specifically Meta’s mbart50-large model [4], was selected to perform fine-tuned summarization of legal texts in Spanish. Built on the BART architecture, this sequence-to-sequence (seq2seq) Transformer model employs an encoder-decoder framework originally optimized for machine translation tasks (Graph 3.1).

Encoder-Decoder Model

Graph 3.1



Source: Authors’ own elaboration

The decision to leverage a translation-oriented model, rather than one trained exclusively for summarization, was based on the structural similarity between the two tasks: both involve extracting semantic content from a source text and generating

³ Different minimum length thresholds were tested by randomly selecting text chunks ranging from 300 to 1,500 characters. Human annotators found that a minimum of 600 characters struck a balance between context and content.

coherent, linguistically accurate output. The model’s multilingual capacity—achieved through pre-training on a broad corpus of languages including Spanish—offers robust linguistic representations crucial for accurately capturing the legal and syntactic complexity of legal documents.

Pre-training for this model involved a denoising autoencoding objective, where the system learns to reconstruct original sentences from corrupted inputs (e.g., masked tokens or shuffled segments). This objective encourages the acquisition of deep contextual embeddings and an understanding of syntax, which is particularly valuable for summarizing dense, domain-specific language. Importantly, this strategy reduces the amount of annotated data required for effective fine-tuning—a critical advantage given the limited size of the training dataset in this application.

Ultimately, the model was chosen for its strong performance in multilingual contexts, its ability to generalize from small datasets, and its effectiveness in capturing the linguistic and semantic features of legal Spanish. Fine-tuning on a curated set of legal text chunks and corresponding expert-generated summaries allows us to exploit these strengths for accurate and context-aware summarization.

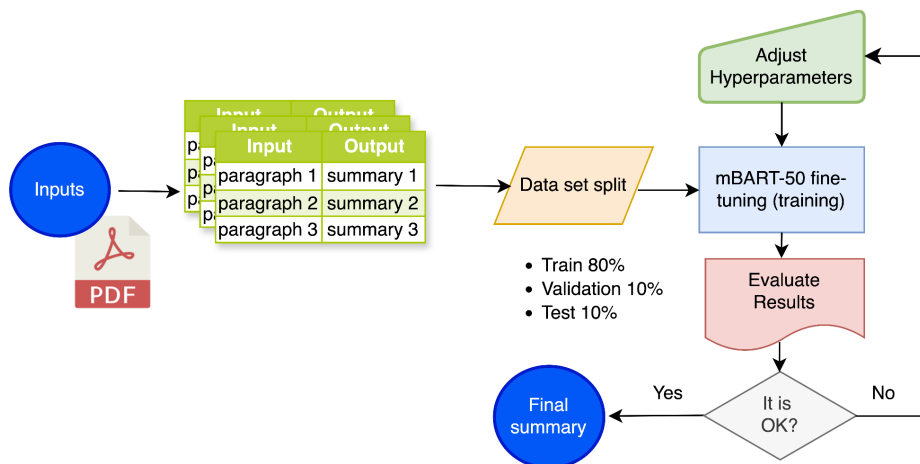
Fine-Tuning Strategy

During fine-tuning, the model receives the segmented legal text chunks produced during preprocessing and is trained to generate concise summaries that closely mirror human-generated references (Graph 3.2). The dataset was split into training (80%), validation (10%), and test (10%) subsets. Training was conducted over 8 epochs using a learning rate of 2e-5 and a batch size of 4.

Approximately 46.7% of the model’s weights were updated during training, while the rest were kept frozen to stabilize training [5]. This selective fine-tuning approach mitigates overfitting and preserves the robust language representations learned during pre-training. It is particularly effective given the limited volume of annotated data, enabling the model to adapt to task-specific nuances while maintaining general linguistic accuracy.

Implementation

Graph 3.2



Evaluation

Model performance was quantitatively evaluated using ROUGE-1 and ROUGE-L metrics. The ROUGE metric (Recall-Oriented Understudy for Gisting Evaluation) is widely recognized as an automated evaluation tool for summarization tasks. Specifically, ROUGE-1 evaluates performance by measuring unigram overlap, reflecting the inclusion and relevance of key terms and essential concepts. ROUGE-L, on the other hand, assesses the structural coherence and syntactic similarity by identifying the longest common subsequence between generated and reference summaries, emphasizing readability and contextual coherence.

Metrics such as ROUGE-2 and ROUGE-SU were omitted due to their limited incremental value over ROUGE-1 and ROUGE-L in capturing conceptual accuracy and readability, particularly given the lower relevance of exact phrasing and skip-bigram matches.

Loss metrics recorded during training and validation phases provided complementary insights into the model's learning behaviour, aiding in the early detection and mitigation of overfitting or underfitting and supporting optimal fine-tuning outcomes.

4. Results

The best-performing model checkpoint was obtained at epoch 4, where training loss reached its lowest point and validation loss remained stable (Graph 4.1), indicating good generalization and minimal overfitting. At this stage, the model achieved a ROUGE-1 score of approximately 35% and a ROUGE-L score of 25% on the validation set (Graph 4.2). These results suggest that the model was able to reliably capture key terms and preserve a moderate degree of structural coherence in its summaries compared to the original reference, despite the complexity and length of the source legal documents.


While a ROUGE-1 score of 35% may appear modest in isolation, it aligns with the range of results reported in prior legal summarization studies, as surveyed by Steffes et al. [6]. For example, Polsley et al.'s CaseSummarizer achieved a ROUGE-1 score of approximately 19%, while Zhong et al.'s iterative masking algorithm reached around 27%. More advanced methods, such as those by Furniturewala et al. and Yoon et al., report higher ROUGE-1 scores ($\approx 49\%$ and $\approx 70\%$, respectively), though these typically rely on extensive training data. As Steffes et al. emphasize, ROUGE scores—while widely adopted—correlate only partially with human judgment in legal domains, where precision, structure, and argumentative clarity often outweigh simple lexical overlap. Within this context, our results reflect a competitive performance given the dataset size and domain complexity.

This trade-off between generality and domain precision is further illustrated in Graph 4.3, which compares summaries generated by ChatGPT (GPT-4o) and our fine-

tuned model for a representative legal text chunk. While ChatGPT produces a fluent summary, it introduces elaborations and paraphrases that risk blurring legal definitions and market distinctions central to the case. In contrast, the fine-tuned model offers a more targeted, compressed rendering that mirrors the legal framing of the original text—demonstrating its alignment with the structural and argumentative style of TDLC proceedings. This qualitative difference underscores the value of fine-tuning over zero-shot prompting for tasks that require legal precision and domain adherence.

Training and validation metrics

Graph 4



To operationalize the model during inference, we adopt a two-stage summarization process that mirrors the structure of the training pipeline. In the first stage, individual summaries are generated for each preprocessed text chunk, allowing the model to handle long documents by breaking them into manageable, semantically coherent segments. In the second stage, these chunk-level summaries are concatenated and passed once more through the model to generate a final, global summary. In addition, a set of descriptive statistics (such as entity count and bigram frequency) is computed to provide contextual information about the case. This hierarchical approach supports coherence at both the local and document-wide levels, ensuring that key legal arguments are preserved while reducing redundancy across sections (Graph 4.4).

Chunk summary example

Graph 4.3

Original Chunk	22.- The advantage of cards associated with international brands - such as VISA- lies in the fact that they are accepted by the entire universe of merchants subscribed to the brand, regardless of the issuer, by virtue of the rule known as \"Honor All Cards.\" This contrasts with cards tied to a proprietary brand, which do not enjoy universal acceptance and whose use is restricted exclusively to merchants affiliated with that proprietary brand.⁵⁰ Consequently, the relevant market affected by the anticompetitive conduct carried out or entered into by VISA, Transbank, and Klap comprises the market of transactions conducted at electronic merchants using credit, debit, and prepaid cards of universal acceptance in Chile, since these types of cards enable cardholders to connect with the e-commerce platforms associated with dlocal, such that proprietary-
----------------	---

	brand cards are not capable of serving as substitutes in this type of transaction.
ChatGPT - GPT 4o	International brand cards like VISA are accepted at all merchants affiliated with that brand, regardless of which bank issues them, thanks to the "Honor All Cards" rule. In contrast, proprietary brand cards can only be used at merchants that accept that specific brand. This means that the market affected by anticompetitive practices by VISA, Transbank, and Klap in Chile is the market for transactions at electronic merchants using universally accepted credit, debit, and prepaid cards. These cards allow users to connect to platforms like dLocal, while proprietary brand cards cannot serve as substitutes in these types of transactions.
Fine-tuned model	The market affected by anticompetitive practices by Visa, Transbank, and Klap is the market for transactions at electronic merchants using universally accepted cards, which enable cardholders to connect to the commerce platforms associated with dlocal.
* Translated from Spanish	

Resumen

Demanda

Ebanx Chile demandó a Visa International por supuestas prácticas anticompetitivas, acusándola de abuso de posición dominante en el mercado chileno de medios de pago. Ebanx, que facilita pagos internacionales sin tarjetas internacionales, alega que Visa busca excluirlo del mercado mediante reglas restrictivas que limitan su acceso a la red de tarjetahabientes. Esto afecta gravemente su modelo de negocio, encarece las transacciones y restringe opciones para consumidores y comercios.

Respuesta a la demanda

Visa International responde a la demanda de Ebanx Chile, argumentando que la Regla de Adquiere Transfronteriza no constituye abuso de posición dominante, sino que asegura la seguridad y funcionalidad de la cadena de pagos, cumpliendo con normativas regulatorias, negando competir directamente con Ebanx, señalando que operan en mercados distintos.

Además refuta las acusaciones de exclusión de mercado, alegando que Ebanx ha presentado información incompleta, incumplido contratos de licenciamiento y distorsionado hechos. Sumado al hecho que su regla se aplica uniformemente a nivel global y garantiza prevención de fraude y cumplimiento normativo.

Participantes

- **Ebanx Chile:** Empresa demandante, agente de recaudo internacional.
- **Visa International:** Empresa demandada, acusada de prácticas anticompetitivas.
- **Transbank S.A., ISwitch S.A., Klap:** Adquirentes locales implicados en la implementación de las reglas de Visa.
- **PayU:** Competidor de Ebanx en pagos internacionales.

Highlights



5. Concluding remarks

This tool demonstrates the feasibility of developing a specialized summarization model tailored to legal documents within a narrowly defined domain. By generating concise, targeted summaries of complex legal texts, the model offers practical support for financial regulation monitoring. In particular, it facilitates the early identification of emerging conflicts in the payments industry and can assist in shaping regulatory frameworks that both reflect market dynamics and promote orderly payment system functioning.

Recognizing the inherent subjectivity in both summarization and regulatory interpretation, the model is designed to incorporate targeted preferences and thematic priorities. Through the use of expert-generated training summaries, domain-specific fine-tuning, and adjustable inference parameters, the system allows regulatory departments to tailor outputs to their analytical needs and strategic goals. As such, the model delivers not only technical utility but also institutional adaptability.

Looking ahead, future work will extend the model's application to market inquiries submitted to the Central Bank, which often involve non-public information and require strict confidentiality safeguards. Because the tool can be securely deployed within the BCCh’s internal infrastructure, it eliminates reliance on external data processing and supports full compliance with internal data protection protocols. This secure, on-premises deployment capability makes the model especially well-

suited for analysing sensitive regulatory materials and serves as a robust foundation for broader applications across internal textual datasets.

6. References

- [1] Banco Central de Chile. Informe de Sistemas de Pago (ISiP), 2024
- [2] Corte Suprema de Chile. Sentencias relevantes sobre adquirencia y sistemas de pago digital. 2024.
- [3] Comisión para el Mercado Financiero. Normativa sobre el Sistema de Finanzas Abiertas, 2024.
- [4] Y. Tang, C. Tran, X. Li, P.-J. Chen, N. Goyal, V. Chaudhary, J. Gu, and A. Fan, "Multilingual translation with extensible multilingual pretraining and finetuning," arXiv preprint arXiv:2008.00401, 2020.
- [5] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, et al., "Huggingface's transformers: State-of-the-art natural language processing," arXiv preprint arXiv:1910.03771, 2019.
- [6] B. Steffes, P. Rataj, L. Burger, and L. Roth, "On evaluating legal summaries with ROUGE," in *Proc. 19th Int. Conf. Artif. Intell. Law (ICAIL)*, Braga, Portugal, Jun. 2023, pp. 457–461.



AI-Driven Legal Summarization for Regulation Monitoring

Central Bank of Chile

4th IFC Workshop on Data Science in Central Banking

Authors: Isidora Díaz J. & Diana López A.

Wednesday 19th February, 2025

Disclaimer


The views are those of the authors and do not necessarily reflect the views of the Central Bank of Chile or its board members.

Agenda

- 1 Context and motivation
- 2 Data
- 3 Model and implementation details
- 4 Results
- 5 Final comments

- The Chilean payment systems industry has experienced significant growth in recent years, particularly in digital payments, which continue to show rapid increases in transaction volume, with new participants entering areas such as payment issuance and acquiring [1]

Figure: Digital Payment Transactions [2]



Note: Credit and debit card data include both businesses and individuals. TEF (Electronic Funds Transfers) data only account for individuals and include both intrabank and interbank transactions.
Source: Central Bank of Chile based on data from the Financial Market Commission (CMF).

- The consolidation of this industry's development has created frictions among market participants, particularly concerning fees and the enablement of new business models [1].
- This has increasingly led some entities to seek resolution from institutions like the Tribunal for the Defense of Free Competition (TDLC) or the National Economic Prosecutor's Office (FNE), which provide legal solutions to these conflicts.
- These dynamics pose regulatory challenges for the Central Bank of Chile (CBCh), as one of its objectives under the Organic Constitutional Law (LOC) is to “ensure the proper functioning of domestic and international payments.”
- Consequently, it is **essential to understand and monitor the cases brought to the TDLC by various stakeholders in the payment systems industry.**

■ Problem:

- Legal documents are lengthy, complex, and use terminology that is not accessible to non-experts, hence the need to have a comprehensive and timely overview of each case.
- Relevant information (issues, entities, and context) is often buried within extensive legal texts.

■ What we do:

- 1 Customize a Large Language Model (LLM) to process original documents.
- 2 Generate one-page summaries highlighting key content in an accessible, high-level manner.

- Includes TDLC cases from 2016 to 2022 related to payment systems.
- Covers both contentious and non-contentious cases (19 in total, avg. 4 documents per case, 40-60 pages per document):
 - Contentious: Judicial disputes with judge intervention.
 - Non-Contentious: Petitioner-only cases; judge intervenes if required by law.
- Analyzed documents that focused on main claims and responses (counterclaims), which vary based on the number of entities involved.
- Each document is then split into chunks, with approximately 150 chunks per document

TDLC legal cases

■ Case Files:

- May include numerous documents per case, in diverse formats with extended file lengths.
- New documents are uploaded regularly (as long as the case remains open).

Estado Diario

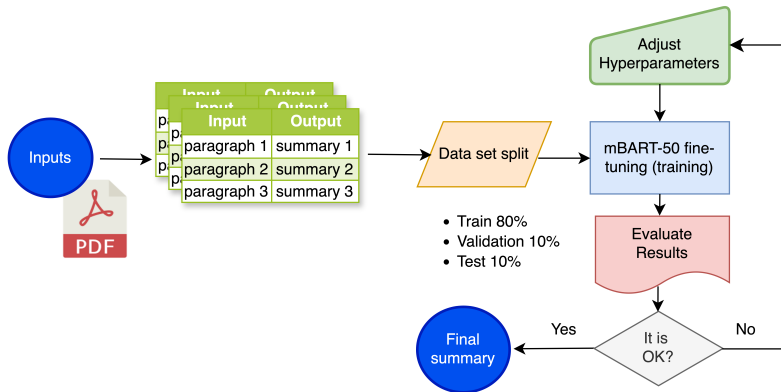
Causa: C-444-2022 D Demanda de Ebanx Chile Ltda. en contra de Servicios Visa Internacional Ltda.
Causa(s) Acumulada(s): C-450-2022 D; C-448-2022 D;

Cuaderno: Cuaderno principal Fecha Ingreso: 06/05/2022

Expediente				Información de Causa		Partes	
Título	Descargar	Detalles	Referencia	Fecha	Parte	Firmantes	Foja
Resolución			Ha lugar	17-12-2024			1151 - 1151
Escrito			Reposición	02-12-2024	DLOCAL CHILE SpA		1150 - 1150
Resolución			Ha lugar	26-11-2024			1149 - 1149
Escrito			Solicita lo que indica	07-06-2024			1113 - 1113
Escrito			Solicita lo que indica	07-06-2024	DLOCAL CHILE SpA		1114 - 1114
Escrito			Solicitud que indica	07-06-2024	Ebanx Chile Limitada		1118 - 1118
Escrito			Solicitud que indica	07-06-2024	Ebanx Chile Limitada		1116 - 1116
Escrito			Solicitud que indica	07-06-2024	Servicios Visa Internacional Limitada		1122 - 1122
Escrito			Solicita lo que indica	07-06-2024	PAVU CHILE S.A.		1120 - 1120
Actuación			Certificación_notifica resolución que indica	15-11-2024			1148 - 1148
Resolución			Se fija nueva fecha de audiencia de conciliación	14-11-2024			1147 - 1147
Escrito			Solicita lo que indica	12-11-2024			1146 - 1146
Actuación			Notifica resolución que indica	05-09-2024			1145 - 1145
Resolución			Suspensión del procedimiento de común acuerdo	03-09-2024			1144 - 1144
Escrito			Solicita suspensión de procedimiento	30-08-2024			1142 - 1142
Resolución			Téngase presente	02-09-2024			1143 - 1143
Escrito			Delega poder	28-08-2024	DLOCAL CHILE SpA		1140 - 1140

Figure: Example of TDLC case and platform

Implementation



- Fine-tuning adapts a pre-trained language model to a specific task (such as summarization) using a smaller, task-specific dataset.
- It updates the model's weights to optimize performance on the new task while leveraging knowledge from pre-training [3].

Implementation details

- Multilingual transformer model: [facebook/mbart50-large](#) [4]
- Input: Legal text chunks.
- Target: Human-made summaries.
- Dataset split: 80% training , 10% validation, 10% test
- Training Configuration:
 - Learning rate: 2e-5
 - Batch size: 4
 - Epochs: 8
 - Trainable parameters: 46.7%
- Evaluation: [loss & ROUGE](#)
- During inference, summaries are generated in chunks to simplify large documents into manageable extracts. A global summary is then produced in 2 stages, by aggregating these chunks and prompting the model again.

Results: Metrics




Figure: Loss over Epochs




Figure: Validation ROUGE Scores (F1)

Results: Example (1/2)

Original chunk

22.- The advantage of cards associated with international brands -such as VISA- lies in the fact that they are accepted by the entire universe of merchants subscribed to the brand, regardless of the issuer, by virtue of the rule known as "Honor All Cards." This contrasts with cards tied to a proprietary brand, which do not enjoy universal acceptance and whose use is restricted exclusively to merchants affiliated with that proprietary brand.⁵⁰ Consequently, the relevant market affected by the anticompetitive conduct carried out or entered into by VISA, Transbank, and Klap comprises the market of transactions conducted at electronic merchants using credit, debit, and prepaid cards of universal acceptance in Chile, since these types of cards enable cardholders to connect with the e-commerce platforms associated with local, such that proprietary-brand cards are not capable of serving as substitutes in this type of transaction.

Results: Example (1/2)

ChatGPT - GPT 4o

International brand cards like VISA are accepted at all merchants affiliated with that brand, regardless of which bank issues them, thanks to the "Honor All Cards" rule. In contrast, proprietary brand cards can only be used at merchants that accept that specific brand. This means that the market affected by anticompetitive practices by VISA, Transbank, and Klap in Chile is the market for transactions at electronic merchants using universally accepted credit, debit, and prepaid cards. These cards allow users to connect to platforms like dLocal, while proprietary brand cards cannot serve as substitutes in these types of transactions.

Fine-tuned model

The market affected by anticompetitive practices by Visa, Transbank, and Klap is the market for transactions at electronic merchants using universally accepted cards, which enable cardholders to connect to the commerce platforms associated with dlocal.

Results: Example

H. TRIBUNAL DE DEFENSA DE LA LIBRE COMPETENCIA

Luis Eduardo Toro Bossay, Francisco Bórquez Electorat, José Luis Corvalán Pérez y Catalina Villalobos Hinojosa, abogados, en representación de Servicios Visa Internacional Limitada, en estos autos caratulados “**Demanda de Ebanx Chile Ltda. en contra de Servicios Visa Internacional Ltda.**”, Rol C N° 444-2022, al H. Tribunal de Defensa de la Libre Competencia (“**H. Tribunal**” o “**H. TDLC**”) respetuosamente decimos:

De conformidad con lo dispuesto en los artículos 20 y 29 del Decreto Ley N° 211 (“**DL 211**”) y el artículo 308 del Código de Procedimiento Civil, contestamos la demanda interpuesta por Ebanx Chile Limitada (“**Ebanx**” o “**Demandante**”), en contra de nuestra representada, solicitando desde ya al H. Tribunal su rechazo en todas sus partes, con costas, en virtud de los antecedentes que pasamos a exponer.

En esta presentación, se verá con claridad la serie de imprecisiones y contradicciones de la demanda que Ebanx ha interpuesto en contra de nuestra representada, en donde ha intentado argumentar que la Regla de Restricción en Adquiere Transfronteriza de Visa (“**Regla de Adquiere Transfronteriza**” o la “**Regla**”) de Visa constituye un abuso de posición dominante, cuando en la práctica, ello no puede distar más de la realidad.

En efecto, la Demandante ha acompañado información incompleta para efectos de crear segmentaciones artificiosas del mercado relevante y confusiones acerca del verdadero efecto de

Resumen

Demanda

Ebanx Chile demandó a Visa Internacional por supuestas prácticas anticompetitivas, acusándola de abuso de posición dominante en el mercado chileno de medios de pago. Ebanx, que facilita pagos internacionales sin tarjetas internacionales, alega que Visa busca excluirlo del mercado mediante reglas restrictivas que limitan su acceso a la red de tarjetahabientes. Esto afecta gravemente su modelo de negocio, encarece las transacciones y restringe opciones para consumidores y comercios.

Respuesta a la demanda

Visa Internacional responde a la demanda de Ebanx Chile, argumentando que la Regla de Adquiere Transfronteriza no constituye abuso de posición dominante, sino que asegura la seguridad y funcionalidad de la cadena de pagos, cumpliendo con normativas regulatorias, negando competir directamente con Ebanx, señalando que operan en mercados distintos.

Además refuta las acusaciones de exclusión de mercado, alegando que Ebanx ha presentado información incompleta, incumplido contratos de licenciamiento y distorsionado hechos. Sumado al hecho que su regla se aplica uniformemente a nivel global y garantiza prevención de fraude y cumplimiento normativo.

Participantes

- **Ebanx Chile:** Empresa demandante, agente de recaudo internacional.
- **Visa International:** Empresa demandada, acusada de prácticas anticompetitivas.
- **Transbank S.A., ISwitch S.A., Klap:** Adquirentes locales implicados en la implementación de las reglas de Visa.
- **PayU:** Competidor de Ebanx en pagos internacionales.

Highlights



Figure: Original document

Figure: Generated summary
(one-pager)

Final comments


- This application makes it possible to create a model capable of producing specialized summaries on a very specific topic, thus helping our unit identify new business models in the payment industry and develop regulations that align with industry growth.
- There is the potential to use this tool for confidential documents, as it can ensure that the information never leaves CBCh's network and servers.
- Finally, this application is flexible enough to be tailored to the preferences and emphasis of a regulatory department manager, considering the subjective nature involved in training a model for this use case.

Bibliography

- [1] C. B. of Chile, “Payment systems report,” August 2023.
<https://www.bcentral.cl/documents/33528/4387742/Informe-de-Sistemas-de-Pago-agosto-2023.pdf/ec39444d-4623-05ee-aa4c-263cbcd72b11>.
- [2] C. B. of Chile, “Payment systems report,” August 2024.
https://www.bcentral.cl/documents/33528/6165052/ISIP_2024.pdf/880f1337-7904-2386-c815-077cfcb7896b.
- [3] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, *et al.*, “Huggingface’s transformers: State-of-the-art natural language processing,” *arXiv preprint arXiv:1910.03771*, 2019.
- [4] Y. Tang, C. Tran, X. Li, P.-J. Chen, N. Goyal, V. Chaudhary, J. Gu, and A. Fan, “Multilingual translation with extensible multilingual pretraining and finetuning,” *arXiv preprint arXiv:2008.00401*, 2020.

Appendix: Model mBART-50

- mBART-50 is built on BART's sequence-to-sequence (seq2seq) Transformer framework, including an encoder and a decoder built for translation.
- It is pre-trained to reconstruct sentences from noisy inputs (masked or shuffled text), enabling it to learn deep linguistic representations.



Appendix: Model evaluation

- **Loss function:** Cross-entropy loss, matching predicted and actual token sequences.
- **ROUGE** (Recall-Oriented Understudy for Gisting Evaluation): used to evaluate the quality of generated summaries, by measuring n-gram overlap between generated and reference summaries.
- However, note that ROUGE ignores semantic meaning and focuses only on surface-level word overlap.
- Human feedback is still valuable to improve the model, as summary quality is subjective and context-dependent.

[Back](#)