Office, Rank, and Mobility: Decomposing Historical Chinese Bureaucratic Titles with Language Models*

Yue Yu[†1], Jun Chen[2], Michael Yan Hon Chung[1], and Cameron Campbell[1]

[1]Hong Kong University of Science and Technology
[2]Renmin University of China

This version: October 21, 2026

This work was supported by HK RGC Areas of Excellence AoE/B-704/22-R (Chen Zhiwu PI).

*We are grateful for discussions and comments from Michael Hout, Siwei Cheng, Wenhao Jiang, Heng Hu, James Evans, Jeffrey Tharsen, Pierre Landry, Guoer Liu, Dongshu Liu, Dong Zhang, Donghyeok Choi, Yueran Hou, and Zheng Zhang. We have also benefited from suggestions of audiences at the New Horizons for AI Research in the Social Sciences and Humanities Conference at the University of Chicago Hong Kong (January 2026), the Chinese Text as Data Conference at the Chinese University of Hong Kong (May 2026), and the RC28 Summer Meeting at New York University (August 2026). Zijie Zheng provided assistance with data annotation.

[†]yue.yu@connect.ust.hk

Abstract

We introduce a procedure that employs an adapted small language model to decompose compound job titles in bureaucracies and other large organizations into substantively meaningful components using a schema grounded in institutional context. This addresses an issue commonly encountered in historical sources for China and neighboring countries according to which recorded job titles include diverse additional attributes beyond the post, such as location, eligibility, temporary assignments, additional titles, or honors, with no consistent format. This makes it difficult to infer the status of officeholders at scale, because their job titles can't be matched to a canonical list that also specifies ranks or other measures of status. It may be unclear whether a changed job title reflected a promotion, a demotion, or a lateral transfer. We demonstrate the procedure by applying it to personnel records of 19th century officials in the China Government Employee Database–Qing (CGED-Q) Jinshenlu (JSL). In the Qing bureaucracy, compound titles were common for senior officials and purchased and awarded honors, temporary assignments, and locations were routinely recorded in job titles. The model achieves average agreement with manual annotation of 94 percent across fields, outperforming larger general-purpose models. The model works for job titles written in other scripts: we apply it to the *hangul* form of the job titles in publicly available career records of Joseon dynasty Korean officials and show that it has high accuracy in identification of departments and posts. Our decomposition thus enables career mobility researchers not only to extract names of posts that can be mapped to administrative rank, but also to separate and make use of other components of job titles relevant to their interests. As a case study, we show that analysis of posts, additional titles, and eligibility extracted from job titles of late Qing officials yields novel insights into dynamics of bureaucratic careers and the composition of officialdom. An analysis of officials from 1830 to 1905 reveals divergent trends after the Taiping Rebellion (1850–1864). While promotion to a higher office within three years declined from 12% to 8%, eligibility for higher office more than tripled, and the share receiving a higher additional title nearly doubled.

1 Introduction

Studies of career mobility and occupational attainment typically rely on a job title for information about a person's rank, income, education, or other status. In longitudinal sources like censuses, surveys, personnel files, and historical registers, changes in the job title may reflect advancement, demotion, or lateral transfer. Differences in the job titles of a parent and child provide evidence of social mobility. Because of the importance of job titles for studying mobility, a large methodological literature has developed that offers approaches for mapping job titles to codes or scales that reflect social or economic standing, and can be meaningfully compared across places and times (Hauser and Warren, 1997; van Leeuwen et al., 2002). Such work generally treats a recorded job title as a noisy description of a single, specific post, with the goal being to infer the post from the job title. Recently, large and small language models have made such mapping faster and more accurate, whether by standardizing historical descriptions or by coding originally open-ended entries in sources (Dahl et al., 2024; Laughlin et al., 2024; Sturgis et al., 2026).

While mobility in organizations like the United States federal civil service, which ranks or grades posts on a scale, can be measured directly as a change of rank (DiPrete and Soule, 1988), in many other cases, there are widely recognized complications. Job titles vary considerably within detailed occupations (Martin-Caughey, 2021). In online job postings, these variations account for much of the variation in posted wages (Marinescu and Wolthoff, 2020). Organizations also use titles to confer recognition and distinguish between employees who do superficially similar work, so that the number and arrangement of titles evolves according to incentives and status competition as well as to the division of labor (Baron and Bielby, 1986; Strang and Baron, 1990; Grant et al., 2014).

In many historical settings, there is another complication: a job title may record auxiliary information about rank, eligibility for appointment to a specified office, location, temporary duties, and honors (Yan, 2010, 327–336). In such scenarios, the recorded job title is a concatenation of the post itself with various attributes of the post or its holder. In these situations, inferring status and studying mobility requires decomposing the various components embedded within a job title to extract the post that is the basis of the assignment of rank or grade, and parsing other attributes that may be relevant for understanding aspects of status other than bureaucratic rank. Since manual decomposition is impractical when the volume of records is large, a coding rule must be applied to extract the substantive post that can be mapped to a rank, and parse the remaining items into separate fields. Because of the diversity and open-ended nature of the job titles recorded for officials, especially senior ones, a keyword-based coding rule that is generally applicable is difficult or impossible to specify, whether by regular expressions or other approaches. In such situations, the presence of additional attributes may lead the inferred rank or grade produced by a keyword-based coding rule to be inaccurate. At the same time, such additional attributes may be of substantive interest in their own right, and a keyword-based coding rule may miss them.

An example showcases the difficulty of a keyword-based approach via regular expressions or other methods, and the necessity of a more flexible approach, which we will suggest should be based on language models. Consider the governor's title in Figure 1: "Vice Minister of War" denotes a

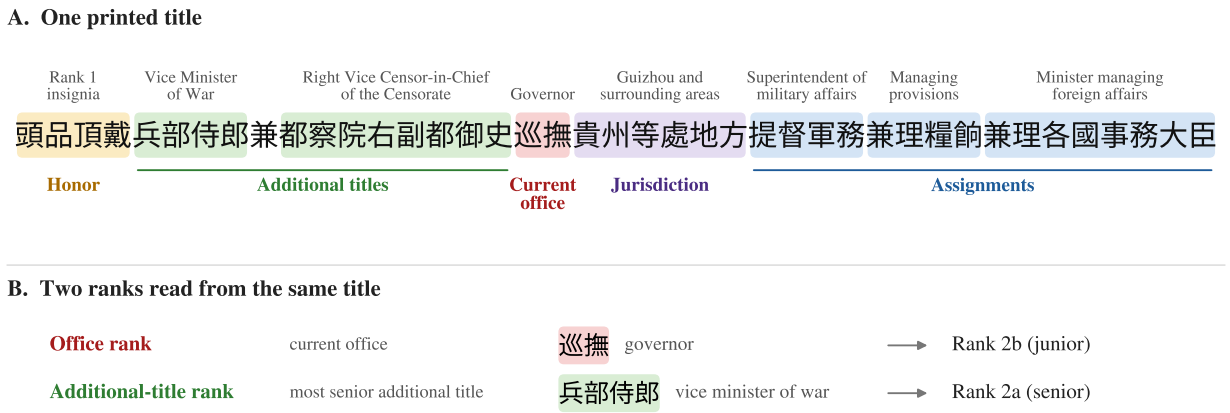

Figure 1: A governor's title recorded in CGED-Q JSL, shaded by institutional component and glossed in English. The current office carries rank 2b, whereas the more senior of the additional titles carries rank 2a. Rank suffixes a and b denote senior and junior grades, respectively. The same notation is used throughout the paper.

central government post, yet here it is actually an additional title attached to a governorship. Elsewhere, the same name might signify an actual post or eligibility to be appointed to such a post, and only its meaning can be inferred from its context, which determines its role. Keyword matching through regular expressions would fail to account for this. A language model, however, is able to interpret each component in relation to the entire job title. Training it on manually annotated titles teaches it how a particular bureaucracy writes its titles. This is knowledge that is difficult to formalize as explicit rules. By contrast, fine-tuning a small model on a modest set of annotations can achieve performance competitive with much larger models on researcher-defined classification tasks (Do et al., 2024; Chae and Davidson, 2026). Furthermore, a small open model can be executed locally and reproducibly across an entire corpus, avoiding the costs and other issues associated with proprietary models (Ollion et al., 2024).

This paper develops and evaluates a procedure based on small language models for decomposition of job titles, testing it in a particularly challenging context: historical records of the job titles of Qing dynasty bureaucrats in 19th-century China. These were included in the *Jinshenlu* (Figure 1), a roster of civil and military officials produced by the central government every three months, originally for the reference of the Emperor and his advisors, but in the 19th century widely available for purchase in the form of commercial editions produced by publishers who had gained access to the official government editions. The job titles in this source routinely included not only a post, but bestowed and purchased additional titles, eligibility, temporary assignments, and locations. We make use of a digitized version of the *Jinshenlu*, the China Government Employee Database–Qing Jinshenlu (CGED-Q JSL) (Chen et al., 2020), which we describe below. By nominative record linkage, the CGED-Q JSL tracks individual officials from one quarterly edition to the next for years or decades (Yu et al., 2026; Campbell and Chen, 2022). We use the decomposed posts, eligibility, and bestowed and purchased additional titles produced by our approach for a case

study of promotion in the late Qing, and time trends in the prevalence of additional titles. We show that during and after the Taiping Rebellion (1850–1864), when the Qing state increasingly sold and awarded eligibility for office and additional titles (Kaske, 2007; Zhang, 2022), serving officials more often advanced in these titles, while promotion to a higher office became less frequent. The decline in promotion is consistent with studies showing that officials left office less often, so that fewer posts became vacant (Campbell and Gao, 2024; Campbell et al., 2026), and that the war enlarged the pool of candidates competing for those that did (Yang, 2021).

We propose a three-stage procedure for job title decomposition (Figure 2). First, we design a schema aligned with institutional knowledge to define the title's components, such as the currently held post, additional titles, eligibility for a higher office, temporary assignments, and bestowed and purchased additional titles. Next, we adapt a small language model, that is, fine-tune it on approximately 2,000 annotated titles sampled from CGED-Q JSL, so that it assigns the expressions within each title to these predefined components. Finally, a rule-based lookup converts the chosen component to a rank. The decomposed job titles are output to a new dataset linked to the original one that allows researchers to follow the component most relevant to their research question, rather than leaving that critical choice to a rigid and likely unreliable coding rule.[1]

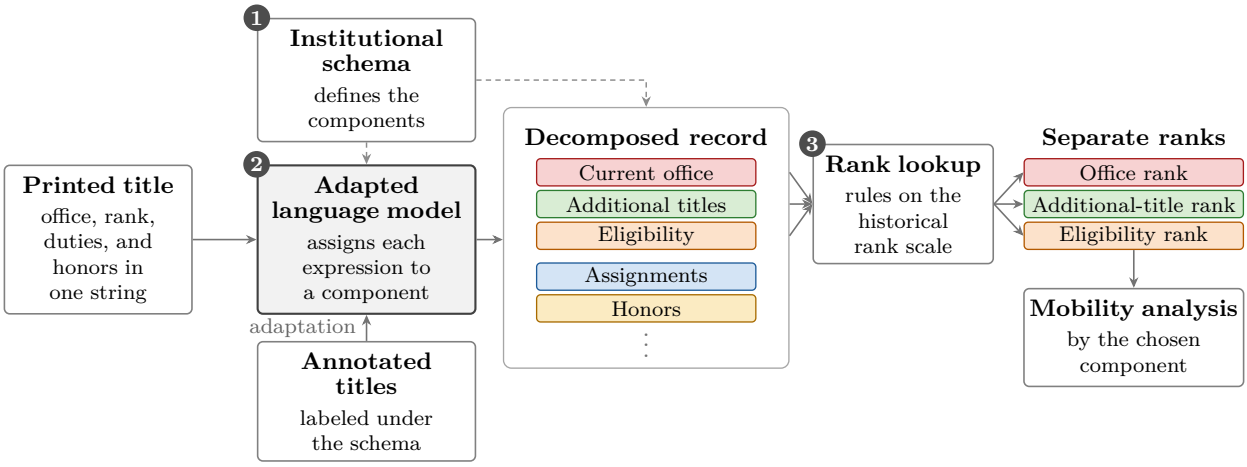

Figure 2: Overview of the procedure. An institutional schema defines the components a title can contain (1). A small language model, adapted on titles annotated under the schema, assigns each expression of a printed title to one of these components (2), and a rule-based lookup converts the rank-bearing components to rank (3). The ranks are kept separate, so that a mobility analysis follows the component its question concerns.

We first compare the results of our approach to job title decomposition with manual annotations of two sets of 200 held-out titles. The adapted model agrees with the annotation on about 94 percent of fields, against 72 to 78 percent for general-purpose models given the same instructions and 72 percent for a rule-based parser built from the same annotations. For most of the components that mobility measures use, the model's reading yields the same rank as the manual annotation. We

[1]We release this dataset for the public version of CGED-Q JSL, together with the adapted model and the scripts of the procedure (Section 8).

then repeat the procedure in a second bureaucracy, using publicly available career records of some 8,700 Joseon Korean officials who passed the civil service examination. With department and post defined for that bureaucracy and the same small model adapted on 1,000 annotated titles, the model reads both components with high accuracy, well above the general-purpose models. Importantly, we use the titles in *hangul* script, showing that our approach is not limited to materials written in Chinese characters.[2]

We showcase the usefulness of our approach for the study of the late Qing bureaucracy by applying it to the CGED-Q JSL rosters from 1830 to 1905, to study trends in promotion and eligibility for a higher office. The decomposed job title allows us to follow 64,000 officials and to measure advancement according to the bureaucratic rank of their post, and as reflected in additional titles, and in eligibility for future promotion. We show that promotion to a higher office became much less frequent in the 1860s. It fell by a third between the decades before and after the war. By contrast, the prevalence of higher additional bestowed and purchased titles increased, as did the recorded eligibility for higher office become more common, so that eligibility and promotion moved in opposite directions.

The paper makes three contributions to the measurement of post from text. First, it reframes compound title coding as a construct choice, advocating decomposition as a foundational design for sources where a single string encodes multiple institutional attributes. Second, it demonstrates the procedure's high adaptability and generalizability by applying it to complex historical contexts (Qing China and Joseon Korea). This validation extends beyond mere extraction accuracy, tracing measurement errors into downstream variables. Third, it reveals that the choice of title component impacts substantive mobility conclusions, thereby requiring clear reporting of the measured component when presenting title-derived mobility rates. Section 2 describes the source data and situates the problem within occupational coding literature. Section 3 identifies the measurement challenges posed by compound titles and presents the decomposition and rank construction procedure. Section 4 presents the evaluation, Section 5 analyzes mobility trends across components, and Section 6 discusses the procedure's broader applicability.

2 Background

2.1 The *Jinshenlu* and Its Titles

The *Jinshenlu* recorded serving and expectant officials in rosters produced every three months. They listed each official under the department and post in which he served. A typical complete edition included about 13,000 civil officials and 8,000 military officers. CGED-Q JSL is a digitized version of these rosters that links the repeated appearances of an official across editions, creating longitudinal career histories (Chen et al., 2020; Campbell and Chen, 2022). We analyze the rosters

[2]The career records are distributed with Choi et al. (2026) at `https://github.com/dhchoi02/JoseonTSI`. They are taken from the index of persons that the National Institute of Korean History publishes with its digital edition of the *Annals of the Joseon Dynasty* (`https://sillok.history.go.kr`), which is released on Korea's Public Data Portal (`https://www.data.go.kr/data/15053645/fileData.do`).

from 1830 to 1905, using the existing person identifiers, and combine the records of each official within a calendar year into one observation. Records that the database flags as invalid, duplicated, or interpolated are excluded. The resulting panel contains 961,799 person-years for 125,605 officials.

Each record includes a job title in which the roster compressed an official's institutional situation into a single string. Figure 1 is a long example, but even short titles commonly combine a post with an eligibility, an assignment, or an additional title. Because the rosters list only men who held posts in the bureaucracy and, in some cases, men assigned to await appointment, titles and eligibility held by men outside the bureaucracy appear only once their holders entered it. The measures below therefore describe rank among officials within the bureaucracy, which suffices for comparing the components in the careers of serving officials, but are not estimates of the prevalence of such titles in the general population.

2.2 Titles as Measures of Authority and Status

Organizational sociology treats job titles as social categories rather than neutral descriptions of work. Titles often proliferate beyond the requirements of the division of labor, acting as "rationalized myths" or "hedonic wages" that decouple ceremony from technical tasks (Meyer and Rowan, 1977; Baron and Bielby, 1986). Strang and Baron (1990) note that such proliferation often reflects status competition rather than functional change. As with educational credentials (Collins, 1979), titles can inflate according to institutional logics rather than skill demand, so that a title-based scale shifts while underlying duties remain unchanged. Because titles carry both bureaucratic authority and symbolic capital (Bourdieu, 2014), simple comparisons across time may fail to distinguish real mobility from title inflation.

The Qing bureaucracy distinguished between authority and status by layering symbolic honors onto the titles of bureaucratic officeholders (Yan, 2010, 256–257). Substantive posts carried statutory ranks, but the state also institutionalized the sale of offices, eligibility, and ranks (*juanna*) as fiscal expedients to manage deficits and as part of elite bargaining (Kaske, 2007; Wu, 2013; Zhang, 2013, 2022). Sales expanded to meet the costs of the Opium War in the 1840s, and during and after the Taiping War these purchased titles and eligibility became increasingly divorced from actual appointment (Kaske, 2007; Zhang, 2022). Existing quantitative research typically simplifies this complexity into a single variable for bureaucratic rank or a binary indicator of whether an official's highest degree was purchased (Ho, 1962; Shi, 2023). By decomposing a job title into its constituent components, we can distinguish substantive administrative promotion as reflected in post-based bureaucratic ranks from the accumulation of purchased or bestowed additional titles.

2.3 From Occupational Coding to Title Decomposition

Research on occupational coding develops methods for converting heterogeneous descriptions in surveys, administrative records, and historical sources into comparable categories. Classification systems such as the Historical International Standard Classification of Occupations (HISCO) provide a common vocabulary for this purpose (van Leeuwen et al., 2002). Automated approaches

include dictionary matching and supervised classifiers, as illustrated by Procode and comparisons of existing autocoders (Savic et al., 2022; Wan et al., 2023). More recent models learn the correspondence between descriptions and codes from annotated examples, both for contemporary survey responses and multilingual historical records (Laughlin et al., 2024; Dahl et al., 2024). In each case, the aim is to give the same code to different descriptions of the same occupation.

A single record can attribute multiple occupations or offices to the same person—a main job and a side job, an inherited household category alongside the work actually performed (Guo et al., 2019), or a current post alongside a former one. These designations are not equivalent, so before any of them can be coded, the researcher must determine how each relates to the person. In surveys, the questionnaire specifies it: the American Community Survey asks about the job held in the previous week (Laughlin et al., 2024), and SOCbot asks follow-up questions when an answer is insufficient (Sturgis et al., 2026). In historical sources, studies infer the relation from the way each designation was recorded. Yang (2022) identifies subsidiary employment in criminal-case memorials from verbs describing what a person was doing at the time of an incident. Dai (2025) codes successive and concurrent occupations listed in lineage genealogies on separate lines. Seo and Jang (2026) recognize official titles in the *Seungjeongwon Ilgi* with a fine-tuned language model. They then use rules keyed to adjacent words to classify each official's relation to an office as appointment, incumbency, former office, or posthumous conferment. In all of these cases, each designation is recorded separately, and its role is signaled explicitly in the original source.

Compound job titles, by contrast, combine several designations in a single string, and their roles are signaled only implicitly. In historical sources for imperial China, an official's title often joins the post with additional titles, eligibility, temporary assignments, locations, or honors. Methods that recognize office names in text, such as keyword matching in MARKUS and named-entity models for job titles (De Weerdt, 2020; Adam et al., 2025), would find several office names in the governor's title from the *Jinshenlu* in Figure 1 but could not determine the role each plays in it. Nor are these roles marked by keywords like those that Seo and Jang (2026) use to classify relations. Some roles follow instead from institutional conventions that are familiar to specialists in the bureaucracy concerned but stated nowhere in the title (Yan, 2010, 328). The Vice Minister of War, for example, is an additional title rather than the post because governors conventionally held this title alongside their own office. Our procedure therefore adapts a language model on annotated titles from the bureaucracy in question so that it learns such conventions from examples, reads each title as a whole, and assigns every expression to a field of a structured record based on the other expressions in the same title.

3 Decomposing Titles and Constructing Ranks

3.1 Measurement Problems

The central challenge is that a job title recorded in a Qing *Jinshenlu* roster, as in many other historical Chinese sources, has no punctuation or other delimiter separating the post from the

expressions around it. The post, additional titles, eligibility, assignments, and honors all appear in one string, not always in the same order, and because the same office names served both as posts and as additional titles, an expression's function cannot be read from its wording alone. In the governor's title in Figure 1, for instance, Vice Minister of War (*bingbu shilang*) and Right Vice Censor-in-Chief (*duchayuan you fu duyushi*) are additional titles attached to the governorship (*xunfu*) rather than posts the governor occupied. Shorter titles pose the same problem. Each of the four titles in Figure 3 places beside the current office a component that is not the post: an eligibility, an assignment, an additional title, or a combination of honor, additional title, and assignments. These components measure different things. The rank of the current office reflects the authority the official exercises, an additional title records the standing the state recognizes in him, and an eligibility indicates an appointment he may receive in the future. Assigning a single rank to the whole string therefore lets the coding rule decide implicitly which component counts as advancement.

A second challenge is that a field must hold the standard name of a component, which often cannot be copied from the title as written. In panel B of Figure 3, the title reads "water works assistant magistrate posted at Nancheng." Recording the whole phrase as the office would give a different code to every assistant magistrate with a different duty or station, whereas the schema records assistant magistrate as the current office, water works as an assignment, and Nancheng as the place of station. In panel C, "rank 4 title" precedes "bureau director," but the rank belongs to a separate additional title rather than to the office: a bureau director held rank 5a, and this official had been granted rank 4 in addition. The values recorded in the fields are therefore normalized institutional names rather than substrings of the title. Interpretation must also resolve character variants, which write the same character in different forms with different Unicode values, and conventional abbreviations, mapping both to the canonical name that the rank lookup requires. It must also tolerate errors in which one character was replaced by a similar-looking one, whether in the production of the rosters or in their digitization.

Our procedure addresses these two problems in three stages (Figure 2). A schema defines the institutional components into which every title is decomposed, an adapted language model assigns the expressions of each title to these components, and a rule-based lookup converts selected components to the Qing rank scale. Ranks are thus assigned to components rather than to the original job titles, and the choice of which components to use is made by the researcher, constructing a measure according to their specific needs during analysis rather than relying on the output of a keyword-based or other rigid coding rule.

3.2 An Institutional Component Schema

We represent each title as a structured record with fifteen fields, grouped in Table 1 into five concepts. Table 5 in Appendix A defines each field and gives its key in the data. The current office is the post that the official holds in the department under which the roster lists him, recorded without rank, status, place, or honors. Some posts, such as that of a provincial education commissioner

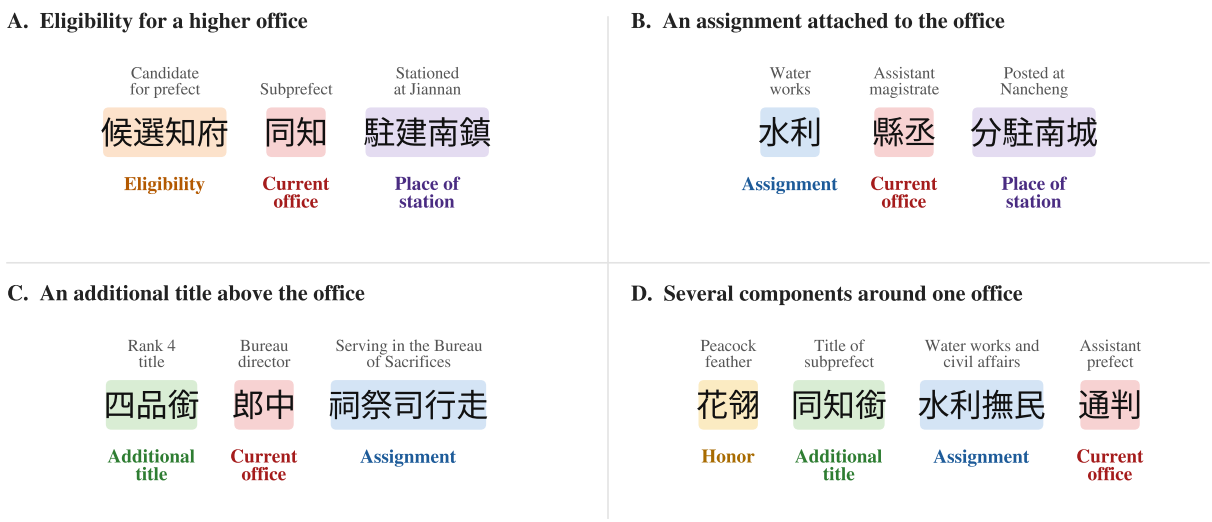

Figure 3: Four titles from the stratified evaluation set, each placing a different component beside the current office. Coding the most senior expression as the post would turn the eligibility in panel A into a promotion and the additional titles in panels C and D into a higher office, whereas the assignments in panels B to D change duties but not the post.

(*xuezheng*), were deputations filled by officials who retained a substantive post elsewhere, and in that case the retained post is recorded as the base office. Two fields record the official's standing with respect to other offices. The office status records the terms on which he holds the current office, for instance as an acting, probationary, or supernumerary holder, and the eligibility records an office to which he is entitled to be appointed but does not yet hold. Assignments are temporary, concurrent, or specialized duties performed alongside the post. Additional titles are office titles or ranks conferred on the official without the duties of the office, such as the vice-ministerial titles held by governors, and the ceremonial office, noble title, martial title, and honors and awards fields record other forms of recognition. The administrative setting is divided into local, central, and military departments, jurisdiction, and place of station. The schema is intentionally more granular than a mobility analysis based on the rank of the post requires, so that researchers can recombine components according to the form of status they study.

Three annotation rules for the manual creation of the training data for the model enforce the distinctions identified above. First, the current office is recorded without rank, status, honor, or geography. Second, the current office must be compatible with the administrative level of its department, and a more senior office expression that cannot be the post exercised in that setting is assigned to an additional-title or eligibility field. Third, the annotation is conservative: a component may be normalized or placed under a broader field, but an office that the title does not express is never supplied, and when a title contains no post the current-office field remains empty.

Table 1: The schema separates the components of a title before any of them is converted to rank. Single-valued fields store one normalized expression, and multi-valued fields store every qualifying expression of their type.

Concept	Fields	Role in measurement
Current post	current office; base office	Authority currently exercised and, for a deputation, the substantive post the official retains
Status and eligibility	office status; eligibility	Terms on which the current office is held, and offices to which the official is entitled to be appointed
Temporary duty	assignment	Temporary, concurrent, or specialized duties, kept apart from the post
Additional rank and honor	additional title; ceremonial office; noble title; martial title; honors and awards	Forms of rank or recognition that may accompany the current office
Administrative setting	local, central, and military department; jurisdiction; place of station	Institutional and geographic location used to interpret the office

3.3 Annotation and Model Adaptation

The annotated collection comprises approximately 2,400 titles. To ensure diverse representation and prevent common titles from dominating the labeling budget, job titles for annotation were specifically drawn from distinct observed strings. The records selected for inclusion in the training data were balanced across time and prioritized strings structurally different from those already chosen. Two Qing history researchers produced annotations according to the established rules. Two independent evaluation sets, each containing 200 titles, were randomly sampled and held out from this collection. The first, a fixed set, was reserved before model adaptation and reflects the collection's overall distribution. The second, a stratified set, was subsequently drawn from the remaining annotations to ensure sufficient representation for scoring all fields, especially rare ones. The remaining approximately 2,000 titles constitute the adaptation pool, with about 1,800 used for training and 200 for checkpoint selection. Thus, a single adapted model is evaluated on both sets, neither of which shares titles with the adaptation pool.

We adapt Qwen3.5-4B with Low-Rank Adaptation (LoRA)(Yang et al., 2025; Hu et al., 2022), a method that trains a small number of additional parameters while keeping pretrained weights immutable. We call the result the adapted model, and the same pretrained model used without this training the unadapted model. The model receives field definitions, institutional constraints, and a single title, subsequently returning a JSON record with a fixed set of keys. Because the annotations record standard names rather than substrings of the title, the model learns to write each component in its standard form as well as to separate it from the others. We employ greedy decoding with zero temperature, ensuring that any given title consistently receives the same deterministic reading.

This procedure does not depend on a particular language model, and in principle models other than Qwen3.5-4B with LoRA could be used. Our procedure is defined by its schema, annotated

titles that reflect annotators' tacit knowledge of title interpretation, and the standard rules converting selected components to rank. While the adapted model reproduces these annotations, any equally accurate extraction system implementing the same components also implements our measure. This decoupling allows independent assessment of two questions: a model's accuracy in recovering annotated components (Section 4), and how component selection alters researchers' quantitative findings (Section 5).

3.4 From Components to Rank

The decomposed components are converted to rank through a rule-based lookup that matches their normalized names to the Qing rank scale. The scale has nine ranks, each divided into a senior and a junior grade, and lower numbers denote more senior ranks. We write the grades as a and b: rank 3a is senior rank 3, and rank 3b is junior rank 3, encoded in computation as 3 and 3.5. The lookup reads an office name together with its department, because the same office could carry different ranks at different administrative levels: a magistrate of an ordinary department held rank 5b, whereas a magistrate of an independent department held rank 5a. The resulting *office rank* is the statutory rank of the current office. When the current office is a deputation without a statutory rank, such as that of a provincial education commissioner, the office rank is that of the base office, the substantive post the official retains.

The same lookup assigns ranks to the other rank-bearing components. An additional title receives the rank of the office or grade it names, and an eligibility the rank of the office for which its holder is eligible, so that a candidate for prefect has the eligibility rank of a prefect. Honors and noble titles retain their text but fall outside the rank scale, and posts below the nine ranks, such as jail wardens, are recorded as unranked (*weiruliu*). A component that the lookup does not contain is left unmatched rather than given the rank of the most similar listed office. Exact matching already places most components on the rank scale (Section 4.2), and the few miswritten names that remain can be checked by hand and added to the lookup as variants, whereas approximate matching would introduce errors of its own, since office names that differ by a single character, such as prefect (*zhifu*) and department magistrate (*zhizhou*), carry different ranks. We likewise keep the ranks of different components separate rather than combining them (for instance, by taking the most senior), because a combined rank would pre-determine which components constitute advancement. Because the decomposed text and the separate ranks are retained, a study can extend the lookup or follow the component most relevant to its question without running the model again, as demonstrated in Section 5.

4 Evaluation

The evaluation of the procedure focuses on three questions that a user of the decomposed titles is likely to need answered. The first is whether the adapted model recovers the components that a manual annotation assigns to unseen titles. The second is whether its remaining disagreements

change the ranks that enter a mobility estimate, since a disagreement over the wording of an office matters less than one over its rank. The third is whether the procedure, rather than the trained model, can be carried to another bureaucracy.

4.1 Agreement with the Manual Annotation

We compare the components that each system returns with the manual annotation on the two evaluation sets described in Section 3.3. Two sets are needed because the components of a title occur with very different frequencies. Almost every title contains a current office, whereas noble titles, ceremonial offices, and martial titles each occur in fewer than 3 percent of titles. The fixed set, which resembles the annotated collection as a whole, therefore shows how well a system reads ordinary titles, but it contains too few of these rare components to score them. The stratified set was drawn so that every field occurs in at least five of its titles, and results reported for individual components use this set.

A field that holds a single expression, such as the current office, is scored as correct when the system's expression agrees with the manual annotation after normalization. A field that can hold several expressions, such as assignments, is scored by F1, which combines two kinds of error: precision, the share of the system's expressions that the annotation also contains, and recall, the share of the annotation's expressions that the system finds. F1 is the harmonic mean of the two, so that it reaches 100 percent only when the system finds every expression and adds none. We summarize these scores in three ways. The field average is the mean score over all fifteen fields and all titles, and measures how accurately the components are read overall. A valid record is an output that can be read as a record containing every field of the schema with a value of the expected type, so that it can enter the data without manual repair. Complete agreement is the share of titles on which all fifteen fields agree with the manual annotation, a strict summary in which a single differing field makes the whole title count as wrong.

We compare our adapted model against four baselines. The first, the unadapted model, is the same model given the same instructions and decoding rule without the training on annotated titles, and shows what the instructions and the pretrained model achieve by themselves. The second is a rule-based parser, which learns the vocabulary of each field from the training annotations and returns the expressions of that vocabulary found in a title, preferring longer matches; it shows how much exact matching of terminology can solve. The third and fourth are two general-purpose large language models, Gemini 3.8 Flash and DeepSeek V4 Pro, which test whether models that have not been adapted to the source recover the same components from these instructions. All language models receive the same input and use temperature zero (more technical specifications in Appendix B).

The adapted model improves both the assignment of components and the regularity of its output (Table 2). Its field average is about 94 percent on both sets and every output is a valid record, whereas the unadapted model yields low field scores and frequently fails to produce one. The rule-based parser is a stronger baseline, reaching about 72 percent, because much of the vocabulary of

Table 2: Agreement with the manual annotation on the two evaluation sets of 200 titles each, whose titles do not occur in the adaptation pool. The field average is the mean score over fifteen fields, using exact agreement for single-valued fields and F1 for multi-valued fields. Values are percentages; the adapted model is shown in bold.

Evaluation set	System	Field average	Valid record	Complete agreement
Fixed	Unadapted model	45.75	10.5	0.0
	Rule-based parser	73.00	100.0	0.0
	Gemini 3.8 Flash	77.98	100.0	3.5
	DeepSeek V4 Pro	75.82	98.5	1.5
	Adapted model	**94.67**	**100.0**	**49.0**
Stratified	Unadapted model	43.59	5.5	0.0
	Rule-based parser	71.83	100.0	0.0
	Gemini 3.8 Flash	76.40	100.0	4.0
	DeepSeek V4 Pro	72.22	98.5	1.5
	Adapted model	**93.78**	**100.0**	**49.0**

titles recurs and can be found by exact matching. It cannot, however, tell what a matched name stands for: the same office name can be the post an official holds, an additional title, or an office for which he is eligible, and correctly decomposing them relies on understanding the context, which a rule-based parser lacks. The general-purpose models almost always return a valid record but reach field averages of only 72 to 78 percent, slightly above the rule-based parser; their records are well formed, but the content of the fields is often wrong. Because instructions are held constant across models, the adapted model's superior performance is driven by the domain knowledge it acquires from annotated examples, capturing subtle patterns in Qing titles that are difficult to formalize through instructions alone.

Neither the field average nor complete agreement adequately captures the accuracy of individual components. Complete agreement, at 49.0 percent on both sets, is low because with fifteen fields to match it is sensitive to any single disagreement. The field average, conversely, is inflated since an empty field counts as agreement: most titles contain no additional title, for instance, so that this field scores highly even though the model reads additional titles less accurately in the titles that do contain one (Table 7). A study should therefore check how accurately the components it uses are read in the titles that contain them, rather than rely on the average for the whole record.

4.2 Consequences of Remaining Errors for Rank

A disagreement over a component does not always change its rank: the model may write an office in a shorter or longer form than the manual annotation and still name an office of the same rank. We therefore rank both readings with the lookup of Section 3.4, which also ranks the records used in the mobility analysis, and compare them (Table 3). On the stratified set, the model's reading receives the same rank as the manual annotation for the large majority of ranked components, and agreement on the current office is similar on the fixed set. The remaining errors differ in kind between components. An error in the current office gives it a different rank about as often as it

leaves it without one, whereas an error in an additional title or eligibility almost always misses the component altogether; the model rarely assigns either of them a wrong rank. These are the ranks that enter the mobility analysis. Because the model reads a given title the same way wherever it appears, each error recurs in every record with that title rather than occurring at random. In a comparison of two years, an advance can therefore be created or hidden when only one of an official's two titles is misread.

Table 3: Ranks of the manual annotation and of the adapted model on the stratified evaluation set. Each row takes the components to which the manual annotation assigns a rank and shows whether the model's reading receives the same rank, a different rank, or no rank. As in the mobility analysis, a deputation takes the rank of its base office.

Component	Ranked in annotation	Same rank	Different rank	No rank from model
Current office	174	158	9	7
Additional title	41	34	0	7
Eligibility	34	28	1	5

Some components receive no rank because the lookup does not contain them, and this happens to the manual annotation almost as often as to the model's reading. A superintendent whose title does not name the post he retains, or a section chief (*kezhang*) in one of the ministries created in the reforms of the 1900s, has no statutory rank in the lookup however the title is read. Of the 43 current offices in the two evaluation sets that the lookup cannot place in the annotation, 37 cannot be placed in the model's reading either. Such titles are rare in the corpus as a whole: only about 2 percent of the offices named in the records are missing from the lookup (Appendix D).

4.3 Transfer to Joseon Career Records

The final evaluation asks whether the procedure can be repeated in another bureaucracy for which the available sources are in a different script. We use the career records of Joseon Korean officials (1392–1910) compiled by Choi et al. (2026), which follow the officials who passed the civil service examination through the appointments recorded in the *Annals of the Joseon Dynasty*. The records give office names in both Chinese characters (*hanja*) and Korean *hangul* script. We use the *hangul* form, together with structured fields for its rank title, department, and post in Korean, taken from the index of persons that the National Institute of Korean History publishes with its digital edition of the Annals. Using instructions written in Korean, the model processes the full string to extract the department and post in the script format used by the structured fields. Because the structured fields are standardized and do not always match the printed title, we keep only titles for which they give the correct reading (Appendix E). From these titles, a stratified draw selects 250 for the evaluation set and 1,000 others for the adaptation pool. The external data support a reliable test of department and post but not of the full Qing schema, and the target is defined accordingly.

We adapt the same base model on the 1,000 titles of the adaptation pool, choosing the learning rate and the number of training epochs by cross-validation within the pool (Appendix E), and

Table 4: Transfer evaluation on 250 Joseon titles that do not occur in the adaptation pool. Every system returns a valid record for every title. The component average is the mean of department and post accuracy. Values are percentages; the adapted model is shown in bold.

System	Both exact	Department	Post	Component average
Unadapted model	34.4	70.0	51.2	60.6
Gemini 3.8 Flash	48.0	63.6	79.6	71.6
DeepSeek V4 Pro	62.0	74.4	84.8	79.6
Adapted model	**95.2**	**97.2**	**98.0**	**97.6**

compare it with the baselines of Section 4.1 except the rule-based parser. The adapted model identifies both department and post correctly in 95.2 percent of titles, against 62.0 percent for the stronger general-purpose model and 34.4 percent for the unadapted model (Table 4). Figure 4 shows four titles from the evaluation set on which the adapted model is correct and the unadapted model is not. Their errors share one mechanism: without adaptation, the model copies words next to the post into the post field, so that an assignment, a mark of concurrent appointment, a status as former holder, or a grade prefix becomes part of the office. As with the Qing titles, adaptation allows the model to internalize the structural logic of Joseon titles by training on examples.

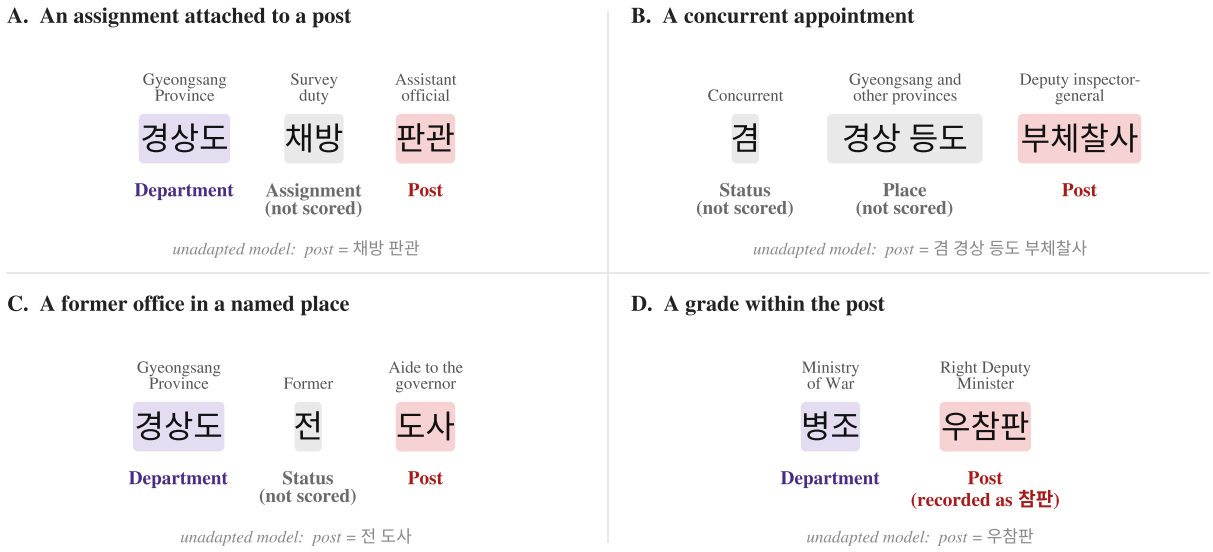

Figure 4: Four Joseon titles from the transfer evaluation set, in *hangul*. The adapted model recovers department and post in each, while components outside the target, shaded gray, remain in the title without being scored. The line beneath each title gives the incorrect reading of the unadapted model.

The transfer test therefore shows that the procedure is not limited to one script or one source. Department and post are concepts that the two bureaucracies share, but the labels, instructions, and examples were defined for the Joseon source, and a modest local annotation sufficed to repeat the sequence of defining institutional components, adapting an extractor, and validating it on

unseen titles.

5 Application: Mobility in Office, Additional Titles, and Eligibility in the Late Qing

The decomposed records make it possible to measure status change separately in three components of the title. A promotion changes the current office, a post in an administrative bureaucracy that carried a fixed rank. An additional title, whether bestowed or purchased, raises the rank that the state formally recognizes for the official without changing his post, and an eligibility places its holder in line for a more senior office that he does not yet hold. The Taiping Rebellion (1850–1864), a civil war that left the state short of both revenue and troops, led it to sell additional titles and eligibility on a large scale and to award them for military service, and these titles became increasingly divorced from the holding of office (Kaske, 2007; Zhang, 2022; Yang, 2021). The three components therefore need not have moved together. We ask how often officials advanced in each of them, and whether the answer differs before and after the war.

For each official, we compare the job titles recorded in two editions of the *Jinshenlu*, three years apart, in records linked by the existing CGED-Q person identifiers. The analysis uses the 449,509 such comparisons in which the official holds a post in one of the nine ranks in both editions (Appendix D). The record pairs cover 64,219 officials, about half of those in the panel, since an official enters the comparison only if he is recorded with a ranked office in two years three years apart. Many officials had careers that lasted for less than three years and therefore cannot be included. An official is *promoted to a higher office* when his office rank is more senior in the later year. He *receives a higher additional title* when the most senior additional title he holds in the later year outranks any of the ones he held three years earlier, where holding none counts as the lowest rank, and he *becomes eligible for a higher office* when the same holds for the rank of the office for which he is eligible. The outcomes are not exclusive, since an official may be promoted and receive an additional purchased or bestowed title in the same interval, and each records the change between the two years rather than every appointment in between.

The three outcomes follow different paths over the century (Figure 5). Promotion was the most common form of advancement throughout, but it became markedly less frequent in the 1860s: 12 percent of officials held a higher office three years later in comparisons beginning before 1864, against 8 percent in those beginning after 1885. The other two components moved in the opposite direction. Over the same span, the percentage who received a higher additional title nearly doubled, from 0.7 to 1.3, and the percentage who became eligible for a higher office more than tripled, from 0.13 to 0.48 (Table 9).

Studies of the Qing civil service suggest two reasons why promotion declined. First, fewer posts became vacant. The number of posts in the rosters changed little over the century, but officials left office less often after the 1850s, and departures became frequent again only after 1900 (Campbell and Gao, 2024). Senior posts, in particular, were increasingly held by older officials who had started

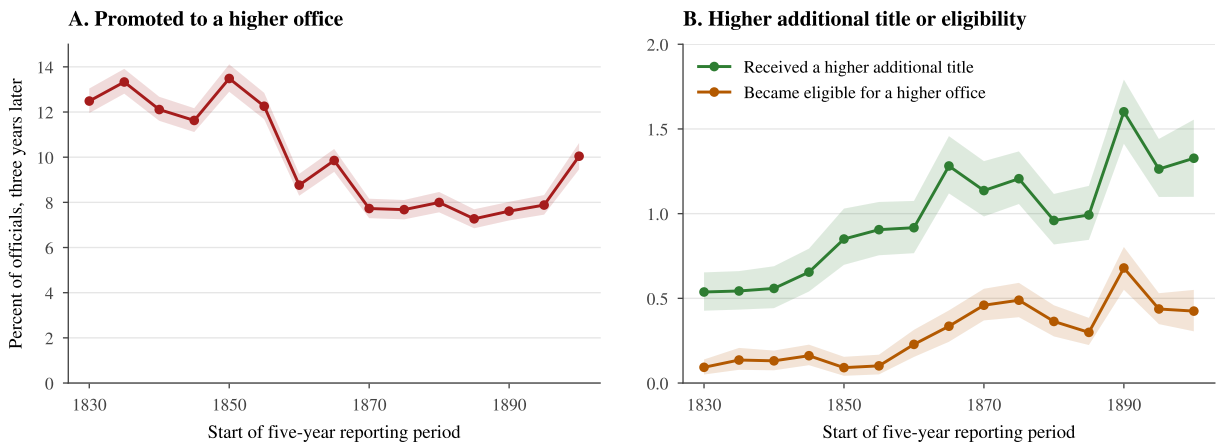

Figure 5: Three-year advancement in each component of the job title. Panel A gives the percentage of officials promoted to a higher office, and panel B the percentages who received a higher additional title or became eligible for a higher office. Each point pools the comparisons that begin in the same five-year period, and shading gives pointwise 95 percent person-bootstrap intervals. The vertical scales differ.

their careers during the Taiping Rebellion and remained in office (Campbell et al., 2026). Because the number of posts changed little, an official could be promoted only when a more senior post fell vacant. Normally, when a vacancy occurred, it could set off a chain of moves as a junior official was promoted into the post, and someone else was promoted into their post, and so on. (White, 1970; Stewman and Konda, 1983). With fewer officials leaving office, fewer posts opened, and there were fewer promotion chains. When departures finally increased in the last years of the Qing, promotion also rose.

Second, more candidates competed for each post that did open. Serving officials due for promotion shared vacancies with candidates who had become eligible through the examinations, purchase, or other routes (Zhang, 2022, 44). The Taiping Rebellion also enlarged the pool of such candidates, since recommendations for military merit and purchase opened new routes to appointment (Yang, 2021). Additional titles and eligibility, by contrast, did not depend on vacancies. An additional title gave no claim to an office, and an eligibility was only a place on a waiting list that the state could extend indefinitely (Zhang, 2022, 33, 36). The number of additional titles and eligibility slots could therefore grow even as the number of actual promotions shrank, as illustrated by Figure 5. Higher additional titles and eligibilities remained far less common than promotion, and their growth did not make up for its decline. They did, however, account for a growing share of all promotions.

Whether officials' chances of advancement shrank after the Taiping Rebellion therefore depends on which components of the title the researcher makes use of. If the components are merged into a single rank, or a single component is selected, such as the most senior rank in the title, their rank would appear to rise whenever that official received a higher additional title. It would count a growing share of advances in title as promotions and make the decline in promotion look smaller than it actually was. It would also miss some real promotions. A county magistrate who already

held the additional title of subprefect, for example, would show no rise in rank when he was actually appointed subprefect. A rank based only on the office would capture the decline in promotion but miss the advances in title entirely. Decomposing the title measures both on the same careers, so that a study of mobility based on titles can state which of them it measures.

6 Discussion

The central design choice is to decompose a title before assigning rank or other status measure, so that the component taken to represent a person's post is chosen by the researcher rather than fixed by a coding rule. Section 5 shows that this choice can decide a substantive conclusion: in the *Jinshenlu*, promotion to a higher office became less frequent after the Taiping Rebellion, whereas advances in additional titles and eligibility became more frequent, so that a measure merging the components records the net of two opposed trends omits the nuanced insights in the data. The evaluation in Section 4 shows, in turn, that the decomposition can be carried out across an entire corpus in high agreement with the manual annotation, and that it can be repeated in a second institution and language. In this section, we first discuss the lessons we learned from developing the procedure, then state its limitations, and finally envision the analyses that can be performed in the future.

6.1 Lessons from Building the Measure

The first lesson highlights the schema's central role, demanding an iterative development process to capture the tacit grammar of bureaucratic titles. We began with a draft schema to guide annotation and model adaptation. Recurring disagreements between the model and annotations were then examined to determine if they reflected a model error or a schema deficiency; if the latter, the schema and affected annotations were revised accordingly. For instance, additional titles and eligibilities were initially combined, as both denoted a more senior office. However, their institutional meanings differ, because an additional title confers a rank without the post, whereas an eligibility places its holder in line for an office he may later hold, and Section 5 shows that the two rose at different rates after the war. Ultimately, the schema stands as a research product itself, embodying substantive judgments about institutional distinctions, and it deserves the same documentation as the model that applies it.

The second lesson demonstrates the efficacy of small, adapted models. Despite their size, these models outperformed larger general-purpose models in aligning with manual annotations for both Qing and Joseon titles. While linguistically competent, large models struggled with source-specific conventions, including schema distinctions (e.g., additional title vs. eligibility), identifying the primary post, and understanding roster abbreviations and ordering. This suggests that the tacit knowledge encoded in annotations (Section 3.3) is more effectively learned through abundant examples than explicit prompt instructions. Moreover, small models offer practical advantages for measurement: they can be run locally over an entire corpus, and with deterministic parameters,

instructions, and results, variations in model output cannot skew analysis findings (Ollion et al., 2024). Thus, for schema-defined measures, investing in the schema and its annotations yields greater returns than pursuing larger models.

6.2 Limitations

The procedure has limitations of three kinds. First, it inherits the weaknesses of the language model that implements it. Such models can produce fluent output that is not supported by their input (Ji et al., 2023), and their reliability as annotators varies with the task and the prompt in ways that should be checked rather than assumed (Lin and Zhang, 2025). In the decomposed titles, these failures would appear as an expression placed in a field the title does not support, or as a component that the title does not contain. The evaluation measures how often such errors occur in held-out titles and how often they change a rank (Section 4.2), but in a corpus of millions of records they cannot be found by inspection, and even a small error rate can bias estimates computed from the extracted components and invalidate their confidence intervals. Design-based supervised learning may help correct this bias by combining the model's readings with expert annotation of a small random sample of records (Maranca et al., 2025).

Second, the procedure's transferability is constrained by its reliance on institutional expertise and substantial annotation labor. Applying it to another bureaucracy requires specialists to define a sophisticated schema and annotate thousands of examples. Such expertise is scarce for most premodern states. The Joseon case avoided this step, because an existing open database already defined the structure of each title, but for the same reason it covers only two components. The Joseon bureaucracy, moreover, followed Chinese models, so the test shows transfer to a culturally similar bureaucracy rather than to one with completely unrelated institutions or script.

Third, an adapted model accurately interprets titles only as long as they adhere to the fixed grammar on which it was trained. While temporal balancing of annotated titles accommodates gradual changes of Qing titles, it cannot handle institutional regime breaks, such as the transition from Ming to Qing titles. Across such significant institutional shifts, a single model would misread titles from periods for which it was not adapted, as meanings and levels could differ between dynasties (Hucker, 1988). Dynasties also differed in where an official's rank came from, whether from the post he held or from ranks he held in his own right (Yan, 2001). A Tang official, for instance, could hold a prestige rank of grade five while serving in a post of grade three, whereas a Qing official's rank came mainly from his office, with additional titles added on top. Consequently, each regime requires a new schema, a new set of annotated titles, and a separately adapted model. Comparing components across such a series necessitates applying the procedure for each regime and establishing correspondence between their respective schemas.

6.3 Further Analyses of the Decomposed Records

Retaining every component in the decomposed records supports analyses that a single rank cannot. For instance, fields like jurisdiction and department allow geographic mobility to be tracked

independently of promotion. This facilitates measuring movements between capital and provinces, among provinces, and the postings of expectant officials without predetermining how location is encoded in a title. Such measures could reveal if promotion increasingly required inter-provincial relocation, or if expectant officials accumulated in specific provinces as eligibility expanded post-war.

The same records also enable measuring title inflation as a property of titles, distinct from individual promotion. Two direct measures emerge from this decomposition. First, the annual distance between the most senior rank acknowledged by a title (via additional titles and eligibility) and the rank of the office held, observed across all officials. Second, the predictive power of holding an additional title or eligibility for later appointment to an office of that rank. Signaling theory posits that a marker loses value when its cost no longer differentiates between those with and without the indicated qualities (Spence, 1973; Collins, 1979). Comparing both measures across decades would reveal whether rank expansion devalued these titles.

6.4 Applications to Other Sources

Historical records in languages that general-purpose models handle poorly due to limitation of specialized knowledge are where adapting a small model is most useful. In the Joseon test, a modest local annotation brought the small model well above general-purpose models many times its size. Other potential applications include the Manchu-language archives of the Qing state itself, a language with few digital resources, and the records of other bureaucracies that wrote compound titles, such as those of premodern Japan and Vietnam. For Chinese sources, local gazetteers and genealogies record titles for many men who never appear in the rosters, including holders of purchased or awarded rank without office. Decomposing these titles would measure how widely rank was held in local society beyond serving officials in the bureaucracy.

Contemporary records present a similar challenge. Write-in responses to occupational questions in censuses and surveys often combine an occupation with employment status, supervisory roles, industry, or employer (e.g., "retired teacher," "owner of a small restaurant"), leading conventional occupational coding to assign a single code (Laughlin et al., 2024; Sturgis et al., 2026). Decomposing these responses into separate attributes empowers researchers to select the most relevant components for their measures, e.g., skills associated with the occupation, or occupation with employment status for class analysis. Similarly, job titles in postings and personnel records combine function with grade, specialty, and organizational unit, exhibiting significant variation even within detailed occupations (Martin-Caughey, 2021; Marinescu and Wolthoff, 2020). For instance, modern Chinese political official résumés often list concurrent party and state posts, with a parallel civil servant grading ladder introduced since 2019, which is challenging for political career studies (Jiang, 2018). Separating grade from function would enable measuring promotion as a change in grade, even in organizations whose records don't explicitly do so (DiPrete and Soule, 1988). It would also allow investigating whether organizations grant more senior titles for unchanged work (Baron and Bielby, 1986; Strang and Baron, 1990).

7 Conclusion

Titles are often the only evidence of where a person stood in a hierarchy, yet many of them record several kinds of post at once. When rank, eligibility, temporary duties, and honors are written into the same string as the post like the Qing bureaucracy in historical China, it is not clear which of them counts as the person's post. This paper has proposed to make it explicit by decomposing each title into its institutional components before any of them is converted to rank. A schema grounded in administrative practice defines the components, a small language model adapted on a set of annotated titles assigns each expression to one of them, and a rule-based lookup converts the chosen component to rank.

The evaluation shows that a small adapted model is well suited to this task. On held-out Qing titles, it agrees with the manual annotation more closely than the unadapted model, a rule-based parser built from the same annotations, and much larger general-purpose models given the same instructions. Its advantage lies in what adaptation teaches: the conventions by which a particular bureaucracy writes its titles, which a dictionary cannot capture and instructions cannot fully state. The same advantage appears in a second bureaucracy and language, where the model adapted on a set of Joseon titles in Korean reads most of them correctly, well above the general-purpose models.

The mobility application shows why the choice of component matters. After the Taiping Rebellion, promotion to a higher office became less frequent among Qing officials, whereas advances in additional titles and eligibility became more frequent, so that a merged rank would have recorded the net of two opposed trends and an office rank only one of them. The same problem arises wherever a single string records a post together with other attributes of standing, as in the personnel records and job postings of contemporary organizations and in write-in responses to occupational questions. Researchers who build mobility measures from such titles should therefore state which component their construct requires and report their results under that choice. Decomposition makes this choice an explicit part of the measure rather than an implicit consequence of the coding rule, and an adapted small model makes the decomposition feasible across an entire corpus.

8 Data Availability and Reproducibility

We have released the decomposed titles along with materials needed to reproduce and extend them. The released decomposed titles cover the same period as the public release of CGED-Q JSL, and are available for download at the same site (Chen et al., 2020) are at `https://doi.org/10.7910/DVN/GMQWVZ`. The adapted model, the scripts, and their documentation are available at `https://github.com/bruceyyu/decompose-chinese-official-titles`.

A Annotation Schema and Model Adaptation

Each decomposed record contains the original title and fifteen fields. Table 5 gives the complete schema, grouped as in Table 1, together with the key under which each field is stored in the data. Single-valued fields return a string or null, and multi-valued fields return an array, which is empty when the title contains no component of that type.

Table 5: Complete title-decomposition schema, grouped as in Table 1. The key is the name of the field in the data.

Field	Key in data	Type	Definition
Current post			
Current office	`duty_office`	String	Post the official currently holds in the department under which he is listed, recorded without rank, status, place, or honors.
Base office	`base_office`	String	Substantive post retained by an official deputed to the current office.
Status and eligibility			
Office status	`office_status`	String	Terms on which the current office is held, such as acting, probationary, or supernumerary.
Eligibility	`qualification_title`	String	Office to which the official is entitled to be appointed but does not yet hold, with any queue position.
Temporary duty			
Assignment	`assignment`	Array	Temporary, concurrent, or specialized duties performed alongside the post.
Additional rank and honor			
Additional title	`additional_title`	Array	Office title or rank conferred on the official without the duties of the office.
Ceremonial office	`ceremonial_office`	Array	Honorary or ceremonial office distinct from current administrative authority.
Noble title	`noble_title`	String	Hereditary or personal noble title.
Martial title	`warrior_title`	Array	Martial title or distinction.
Honors and awards	`honors_awards`	Array	Other decorations, awards, and permissions.
Administrative setting			
Local department	`local_department`	Array	Local administrative units that locate the office within the territorial hierarchy.
Central department	`central_department`	Array	Central-government institutions named in the title.

Field	Key in data	Type	Definition
Military department	`military_department`	Array	Military institutions or commands named in the title.
Jurisdiction	`jurisdiction`	String	Territory or functional domain over which the post exercises authority.
Place of station	`garrison`	String	Stated place of stationing when distinct from jurisdiction.

The annotation rules are relational rather than lexical. A current office must be compatible with the department and administrative level under which the record appears, and an office name that cannot be the post held in that setting is assigned to the additional title or eligibility. The annotator may normalize a known character variant or return a broader institutional label, but never supplies an office that the title does not express. These rules allow the same office name to be read differently when its relation to the rest of the title changes.

The annotated collection contains 2,367 distinct titles, and Table 6 gives the counts for training, checkpoint selection, and evaluation. The fixed set consists of the last 200 titles of the original annotation file. The stratified set was drawn from the remaining titles before training: taking the fields in order from the rarest, it adds titles that contain each field until the field occurs in a fifth of the titles available for it, with a minimum of five and a maximum of thirty, and fills the remaining places at random. Neither evaluation set shares a title with the adaptation pool.

Table 6: Annotation and evaluation partitions.

Partition	Titles	Use
Training	1,770	Adaptation of parameters
Validation	197	Choice of checkpoint by validation loss
Fixed evaluation set	200	Evaluation on ordinary titles
Stratified evaluation set	200	Evaluation of every field, including rare ones

Titles were sampled as distinct normalized strings rather than as roster rows, so that very common offices would not exhaust the annotation budget. The sample was balanced over time and favored strings that lay far apart in a character-based representation of normalized titles, which broadens the range of structures seen in training.

We adapt Qwen3.5-4B with Low-Rank Adaptation of rank 64. Each epoch uses the 1,770 training titles, and the loss is computed on the model's response rather than on the repeated instructions. Empty single-valued fields are returned as null and empty multi-valued fields as empty arrays. Generation uses temperature zero and at most 512 output tokens, and the unadapted model uses the same prompt, prompt template, and decoding settings as the adapted model.

B Baselines and General-Purpose Models

The rule-based parser constructs a separate terminology dictionary for each field from the training annotations alone and matches those expressions as exact substrings of the title. For single-valued fields, it selects the longest match, breaking ties by training frequency and then lexical order. For multi-valued fields, it retains matching expressions unless they are contained within a longer match in that field. The parser does not use the rest of the title, so it cannot tell whether an office name is the post, an eligibility, or an additional title.

The general-purpose models are the production versions of Gemini 3.8 Flash and DeepSeek V4 Pro, evaluated on October 5, 2026. Gemini does not offer a setting without thinking, so we use its lowest setting, whereas DeepSeek receives an explicit setting that disables thinking. Both receive the same prompt as the adapted model, without demonstrations, retrieval, tools, or provider-enforced output formats, at temperature zero. The Qwen models and DeepSeek have an output limit of 512 tokens for Qing titles and 256 for Joseon titles; Gemini has limits of 4,096 and 2,048 tokens because its limit includes reasoning, and no response reaches its limit. The same extraction and scoring code evaluates every model.

C Field-Level Results

Single-valued fields are scored by exact agreement after the normalization used in data preparation, and multi-valued fields by set F1, so that missing and extra components both reduce the score. The field average is the unweighted mean over the fifteen fields. A valid record can be parsed and contains every expected field with a value of compatible type, and complete agreement requires every field of a record to agree.

Table 7 reports the results of the adapted model on the stratified evaluation set, field by field. The overall score covers all 200 titles, including those in which both the manual annotation and the model leave the field empty, whereas the score on positive cases is computed only on the titles in which the manual annotation records the component.

Classifying each disagreement as an omission, an addition, or a mixed boundary disagreement shows where further annotation would help most. Fifteen of the nineteen errors in the base office are omissions, and three of the six martial titles are missed. Errors in assignments, by contrast, more often add an expression, alone or together with an omission, than omit one: 22 are additions only, 23 are mixed, and 7 are omissions only. Table 3 in Section 4.2 shows how often errors in the current office, additional titles, and eligibilities change their rank.

D Panel Construction and Mobility Measures

The longitudinal panel retains the existing CGED-Q person identifiers linked by Yu et al. (2026). Records from the separate military rosters and records flagged as invalid and duplicated in post

Table 7: Results by field on the 200 titles of the stratified evaluation set. Positive cases are the titles whose manual annotation contains the field.

Field	Overall score (%)	Positive cases	Score on positive cases (%)
Current post			
Current office	84.0	198	84.3
Base office	90.5	37	51.4
Status and eligibility			
Office status	97.5	38	89.5
Eligibility	92.0	39	69.2
Temporary duty			
Assignment	77.5	88	65.9
Additional rank and honor			
Additional title	93.0	44	75.0
Ceremonial office	99.5	15	93.3
Noble title	98.0	15	80.0
Martial title	98.5	6	50.0
Honors and awards	97.6	42	91.0
Administrative setting			
Local department	96.9	79	95.9
Central department	96.7	56	89.9
Military department	95.1	65	86.5
Jurisdiction	91.5	38	65.8
Place of station	98.5	41	92.7

are removed. The remaining 2,712,774 records from 1830 to 1905 form 961,799 person-years for 125,605 civil officials. Office rank is the rank of the current office or, when the current office is a deputation without a statutory rank, of the base office. Within a person-year, office rank is taken from the last record of the year that has one, so that a transfer within the year does not produce conflicting observations, and the additional title and eligibility are the most senior printed in any edition of the year. As CGED-Q JSL contains both official and commercial editions of *Jinshenlu* and commercial editions might print fuller titles than official editions (Chen et al., 2020), so that a component printed in one edition may be absent from another. An additional title or eligibility printed in one year is therefore treated as held in later years until the official leaves that office or the office reaches the rank of the component.

Office rank is available for 740,002 of the 961,799 person-years (76.9 percent), including 997 that receive it only through the base office. The lookup first looks up the normalized name of an office as a whole. It also recognizes the variant names listed in its dictionary defined by us, but it never matches a name approximately. Most person-years without office rank (19.5 percent of all person-years) name only posts below the nine ranks: jail wardens account for about half of these, and clerks (*bitieshi*) whose title states no grade for most of the rest. Such clerks are treated as unranked, whereas a clerk whose title states a grade receives it, and because an unranked clerk may hold a grade that his title does not record, comparisons exclude unranked posts altogether. A further 2.6 percent of person-years name an office that the lookup does not contain (2.3 percent of the records that name an office), and 1.0 percent contain no office. The share of person-years with office rank lies between 75.2 and 78.1 percent in every five-year period.

Table 8 gives the ranks that the lookup assigns to each rank-bearing component. Unmatched additional titles and eligibilities are chiefly abbreviated or incomplete office names, expressions that name no office, such as a bare 候補 (expectant), and expressions containing illegible characters.

A comparison joins the same official in calendar years t and $t + 3$ when he holds a post in one of the nine ranks in both. The 449,509 comparisons cover 64,219 officials and 70 observed starting years between 1830 and 1902. Five-year reporting periods are defined by starting years 1830–1834, 1835–1839, and so forth, and the last contains starting years 1900–1902 because the panel ends in 1905. Counts are summed within periods before percentages are taken. Table 9 gives the percentages for starting years before 1864, from 1864 to 1885, and from 1886 to 1902, together with the change between the first and last of these periods, which is computed within each bootstrap replication.

The bootstrap draws 64,219 officials with replacement in each of 1,000 replications, retaining every sampled official's full set of comparisons. The bounds are the 2.5th and 97.5th percentiles, and Figure 5 uses pointwise intervals. The intervals do not reflect uncertainty in the person identifiers or in the readings of the titles.

Table 8: Ranks assigned by the rank lookup to each rank-bearing component in the records from 1830 to 1905, in percent, with the senior and junior grades of each rank combined. The office-rank column counts person-years, as in the mobility analysis, so that its rows add up to the coverage given in the text. In the other two columns, each additional title or eligibility counts once in every record that prints it, so that a record with two additional titles counts twice. The last row gives these totals.

Rank	Office rank	Additional title	Eligibility
Rank 1	0.5	17.1	15.6
Rank 2	1.8	20.5	1.8
Rank 3	0.6	14.8	0.2
Rank 4	2.8	9.6	12.3
Rank 5	8.8	26.1	16.4
Rank 6	5.6	5.4	8.7
Rank 7	15.2	0.5	7.5
Rank 8	30.0	0.2	3.0
Rank 9	11.7	1.2	12.3
Unranked (*weiruliu*)	19.5	0.0	6.8
Unmatched	2.6	4.6	15.5
No office	1.0	–	–
Total counted	961,799	34,554	13,651

Table 9: Three-year advancement by period of starting year. Comparisons gives the number of comparisons in the period with each outcome, and percent divides it by all comparisons in the period (188,076, 144,863, and 116,570). Brackets give 95 percent person-bootstrap intervals; changes are in percentage points.

Starting years	Comparisons	Percent
Promoted to a higher office		
1830–1863	22,681	12.06 [11.80, 12.33]
1864–1885	12,080	8.34 [8.08, 8.60]
1886–1902	9,289	7.97 [7.70, 8.24]
Change, first to last period		−4.09 [−4.46, −3.73]
Received a higher additional title		
1830–1863	1,281	0.68 [0.62, 0.74]
1864–1885	1,631	1.13 [1.04, 1.21]
1886–1902	1,540	1.32 [1.22, 1.42]
Change, first to last period		0.64 [0.52, 0.76]
Became eligible for a higher office		
1830–1863	247	0.13 [0.10, 0.16]
1864–1885	578	0.40 [0.35, 0.45]
1886–1902	558	0.48 [0.42, 0.54]
Change, first to last period		0.35 [0.28, 0.41]

E Construction of the Joseon Transfer Test

The transfer corpus begins with 247,432 career records and 6,305 distinct office strings in the JoseonTSI data (Choi et al., 2026), all of them appointments of officials who passed the civil service examination. Titles are identified by the *hangul* form of the office string, which gives 5,942 distinct titles. The structured fields are standardized and do not always match the title as written, in either direction. A field may supply a component that the title does not contain, as when an office written without its department still receives one, or when an older name of an office receives its later name; conversely, a title may name a department that its field leaves empty. Scoring every record would therefore mix the accuracy of extraction with the way the database standardizes titles.

The benchmark keeps a title only if every record with that title receives the same structured fields, and if the department and post names found in the title are exactly those of the structured fields, with none named where a field is empty. For the retained titles, the structured fields therefore give the correct reading. Because the screen matches the vocabulary of the structured fields against each title, a rule-based parser built from that vocabulary would recover the retained titles by construction, and we do not use it as a baseline. The screen yields 1,284 titles. A deterministic draw, stratified by the components a title contains, the presence of a grade prefix, and the band of the office grade, first assigns 250 titles to the evaluation set and then 1,000 of the remaining titles to the adaptation pool. The learning rate and the number of epochs were chosen by five-fold cross-validation within the adaptation pool, without using the evaluation set, and the model was then trained on all 1,000 titles. Rank titles are not scored because too few examples survive the screen.

The adapted model is trained for four epochs at a learning rate of 5×10^{-4}, and on titles that contain each component it recovers 104 of 107 departments and 241 of 245 posts, or 97.2 and 98.4 percent. The remaining disagreements include boundary choices, such as a modified office name against its shorter head, and apparent omissions in the external database.

F Model Instructions

Each title is sent to a model as a single prompt that consists of the instructions, a request to parse the title, and the title itself between backticks (Tables 10 and 11). The instructions are written in the language of the source, in Chinese for the Qing titles and in Korean for the Joseon titles, and the English translations are not part of the input. The adapted model, the unadapted model, and the general-purpose models receive the same prompt.

Table 10: Instructions for the Qing titles. Each row reproduces one line of the message exactly, next to our English translation, and the last two rows show how the title is appended.

Instructions as given to the model	English translation
角色: 你是一位精通清代官制的歷史學家。你的任務是從包含噪聲（如 NA、-、OCR錯誤）的官職描述中，提取結構化信息。	Role: You are a historian expert in the official system of the Qing. Your task is to extract structured information from descriptions of official titles that contain noise (such as NA, -, or OCR errors).
任務: 請先在內部完成實體邊界切分，再將結果精確映射到指定 JSON 字段。	Task: First segment the title into entities internally, then map the result precisely onto the specified JSON fields.
實體邊界切分時請辨認： 1. 機構/層級（包含地方、中央、武職等）	When segmenting the entities, identify: 1. Institutions and levels (including local, central, and military)
2. 核心官職與本職	2. The core office and the base office
3. 差使/任務（臨時性或專門性職責）	3. Assignments and tasks (temporary or specialized responsibilities)
4. 前後綴/加銜/榮譽/爵位	4. Prefixes and suffixes, additional titles, honors, and noble titles
JSON 字段定義: - original_title: 原始字符串，禁止修改	JSON field definitions: - `original_title`: the original string, which must not be modified
- core_office: 核心實缺/差遣官職，即官員實際行使職權的職位	- `core_office`: the core substantive or commissioned office, that is, the post in which the official actually exercises authority
- base_office: 原隸屬本職（若無差遣為 null）	- `base_office`: the substantive office to which the official originally belongs (null if there is no commission)
- local_department, central_department, military_department: 隸屬機構數組	- `local_department`, `central_department`, `military_department`: arrays of the institutions to which the post belongs
- office_status: 署理、候補等狀態	- `office_status`: a status such as acting or expectant
- qualification_title: 名義資歷（如盡先補用）	- `qualification_title`: a nominal qualification (e.g., *jinxian buyong*, first in line for appointment)
- jurisdiction: 管轄區	- `jurisdiction`: area of jurisdiction
- garrison: 駐地	- `garrison`: place of station
- assignment: 差使任務數組	- `assignment`: array of assignments
- additional_title: 加銜數組	- `additional_title`: array of additional titles
- ceremonial_office: 榮譽虛銜數組	- `ceremonial_office`: array of honorary nominal offices
- noble_title: 爵位	- `noble_title`: noble title
- warrior_title: 勇號數組	- `warrior_title`: array of martial titles
- honors_awards: 封號賞賜數組	- `honors_awards`: array of conferred titles and awards

Instructions as given to the model	English translation
規則: 1. 語法解構與純淨度：提取 core_office 或 base_office 時，必須剝離品階、狀態和榮譽，分別放入適當字段。core_office 必須是純淨的職權名詞。 2. 層級空間約束：core_office 必須與 local_department 的最終行政層級匹配。高於當前機構層級的官名應判為 additional_title 或 qualification_title。 3. 實體守恆：只能修復錯別字或向上歸類，禁止捏造文本中不存在的官職。找不到實質核心職務時，core_office 必須為 null。 4. 核心官職標準化：地理/轄區前綴移入 local_department 或 jurisdiction。遇到「特定前綴 + 通用官職」時，通用部分放入 core_office，特定前綴放入 jurisdiction。多個實缺中選擇與當前部門匹配且權力最大者。	Rules: 1. Grammatical decomposition and purity: when extracting `core_office` or `base_office`, strip rank, status, and honors and place each in the appropriate field. `core_office` must be the bare name of the post that carries authority. 2. Administrative-level constraint: `core_office` must match the final administrative level of `local_department`. An office name above the level of the current institution should be judged `additional_title` or `qualification_title`. 3. Conservation of entities: only correct miswritten characters or classify upward, and never invent an office that does not appear in the text. When no substantive core post can be found, `core_office` must be null. 4. Standardization of the core office: move geographic or jurisdictional prefixes into `local_department` or `jurisdiction`. For a "specific prefix + generic office," put the generic part in `core_office` and the specific prefix in `jurisdiction`. Among several substantive posts, choose the one that matches the current department and carries the most authority.
輸出格式: 只輸出一個合法 JSON object，不要輸出分析、說明或 Markdown code fence。空的單值字段填 null，空的多值字段填 []。	Output format: Output only one valid JSON object, without analysis, explanation, or a Markdown code fence. Fill empty single-valued fields with null and empty multi-valued fields with [].
請解析這個官職： `⟨title⟩`	Please parse this official title: `⟨title⟩`

Table 11: Instructions for the Joseon titles, reproduced line by line with our English translation. The last two rows show how the title is appended.

Instructions as given to the model	English translation
역할: 당신은 조선시대 관직명을 구조화하는 역사 연구자입니다.	Role: You are a historian who structures the office titles of the Joseon period.
과제: 주어진 관직명에 실제로 적혀 있는 구성 요소만 다음 JSON 필드로 분리하십시오. - original_title: 입력 문자열을 그대로 복사 - department: 관청·기관 명칭. 없으면 null - position: 구체적인 직위 명칭. 없으면 null	Task: Separate only the components actually written in the given office title into the following JSON fields. - `original_title`: copy the input string unchanged - `department`: the name of the government office or institution; null if there is none - `position`: the name of the specific position; null if there is none
규칙: 1. 입력에 적혀 있지 않은 기관이나 직위를 추론하여 보충하지 마십시오. 2. 설명이나 Markdown 없이 정확히 하나의 JSON object만 출력하십시오.	Rules: 1. Do not infer and supply an institution or position that is not written in the input. 2. Output exactly one JSON object, without explanation or Markdown.
분석할 관직명: `⟨title⟩`	Office title to analyze: `⟨title⟩`

References

Adam, R., K. Venglarova, and G. Vogeler (2025, Apr). Exploring historical labor markets: Computational approaches to job title extraction. *Journal of Data Mining & Digital Humanities NLP4DH*.

Baron, J. N. and W. T. Bielby (1986). The proliferation of job titles in organizations. *Administrative Science Quarterly 31*(4), 561–586.

Bourdieu, P. (2014). *On the State: Lectures at the Collège de France, 1989–1992*. Cambridge: Polity Press. Edited by Patrick Champagne, Rémi Lenoir, Franck Poupeau, and Marie-Christine Rivière.

Campbell, C. and B. Chen (2022). Nominative linkage of records of officials in the China Government Employee Dataset-Qing (CGED-Q). *Historical Life Course Studies 12*, 233–259.

Campbell, C. and S. Gao (2024). Qingdai wenguan de zuzhi renkouxue yanjiu (1830–1911) [The organizational demography of the Qing civil service (1830–1911)]. *Shehui kexue yanjiu [Social Science Research]* (1), 157–169. In Chinese.

Campbell, C., S. Gao, and C. Li (2026). Qingdai wenguan de nianling dongtai yanjiu (1830–1911) [Research on the age dynamics of Qing civil officials, 1830–1911]. *Jindaishi yanjiu [Modern Chinese History Studies]* (2), 93–116. In Chinese.

Chae, Y. and T. Davidson (2026). Large language models for text classification: From zero-shot learning to instruction-tuning. *Sociological Methods & Research 55*(2), 501–567.

Chen, B., C. Campbell, Y. Ren, and J. Lee (2020). Big data for the study of Qing officialdom: The China Government Employee Database-Qing (CGED-Q). *Journal of Chinese History 4*(2), 431–460.

Choi, D., J. Cha, and J. M. Park (2026). Total Success Index and the longitudinal dynamics of bureaucratic stratification in Joseon Korea. *Physica A: Statistical Mechanics and its Applications 688*, 131353.

Collins, R. (1979). *The Credential Society: An Historical Sociology of Education and Stratification*. New York: Academic Press.

Dahl, C. M., T. Johansen, and C. Vedel (2024). Breaking the HISCO barrier: Automatic occupational standardization with OccCANINE.

Dai, Y. (2025). Lineage genealogies as a new source for researching the occupational structure of twentieth-century China: Tradition (partially) transformed. *Historical Methods: A Journal of Quantitative and Interdisciplinary History 58*(1), 54–79.

De Weerdt, H. (2020). Creating, linking, and analyzing Chinese and Korean datasets: Digital text annotation in MARKUS and COMPARATIVUS. *Journal of Chinese History 4*(1), 17–42.

DiPrete, T. A. and W. T. Soule (1988). Gender and promotion in segmented job ladder systems. *American Sociological Review 53*(1), 26–40.

Do, S., É. Ollion, and R. Shen (2024). The augmented social scientist: Using sequential transfer learning to annotate millions of texts with human-level accuracy. *Sociological Methods & Research 53*(3), 1167–1200.

Grant, A. M., J. M. Berg, and D. M. Cable (2014). Job titles as identity badges: How self-reflective titles can reduce emotional exhaustion. *Academy of Management Journal 57*(4), 1201–1225.

Guo, Y., Z. Zhang, B. van Leeuwen, and Y. Xu (2019). A view of the occupational structure in Imperial and Republican China (1640–1952). *Australian Economic History Review 59*(2), 134–158.

Hauser, R. M. and J. R. Warren (1997). Socioeconomic indexes for occupations: A review, update, and critique. *Sociological Methodology 27*(1), 177–298.

Ho, P.-T. (1962). *The Ladder of Success in Imperial China: Aspects of Social Mobility, 1368–1911*. New York: Columbia University Press.

Hu, E. J., Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen (2022). LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.

Hucker, C. (1988). *A Dictionary of Official Titles in Imperial China*. Southern Materials Center.

Ji, Z., N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, and P. Fung (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys 55*(12), 1–38.

Jiang, J. (2018). Making bureaucracy work: Patronage networks, performance incentives, and economic development in China. *American Journal of Political Science 62*(4), 982–999.

Kaske, E. (2007). The price of an office: Venality, the individual and the state in 19th-century China. In R. J. Antony and J. K. Leonard (Eds.), *Dragons, Tigers, and Dogs: Qing Crisis Management and the Boundaries of State Power in Late Imperial China*, pp. 278–304. Ithaca: Cornell University Press.

Laughlin, L., X. Song, M. Wisniewski, and J. Xu (2024, April). From job descriptions to occupations: Using neural language models to code job data. Unpublished manuscript, U.S. Census Bureau, University of Pennsylvania, and Pennsylvania State University.

Lin, H. and Y. Zhang (2025). Navigating the risks of using large language models for text annotation in social science research. *Social Science Computer Review 0*(0), 08944393251366243.

Maranca, A. R. P., J. Chung, M. Hinck, A. D. Wolsky, N. Egami, and B. M. Stewart (2025). Correcting the measurement errors of AI-assisted labeling in image analysis using design-based supervised learning. *Sociological Methods & Research 54*(3), 984–1016.

Marinescu, I. and R. Wolthoff (2020). Opening the black box of the matching function: The power of words. *Journal of Labor Economics 38*(2), 535–568.

Martin-Caughey, A. (2021). What's in an occupation? Investigating within-occupation variation and gender segregation using job titles and task descriptions. *American Sociological Review 86*(5), 960–999.

Meyer, J. W. and B. Rowan (1977). Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology 83*(2), 340–363.

Ollion, É., R. Shen, A. Macanovic, and A. Chatelain (2024). The dangers of using proprietary LLMs for research. *Nature Machine Intelligence 6*(1), 4–5.

Savic, N., N. Bovio, F. Gilbert, J. Paz, and I. Guseva Canu (2022, 01). Procode: A machine-learning tool to support (re-)coding of free-texts of occupations and industries. *Annals of Work Exposures and Health 66*(1), 113–118.

Seo, M.-J. and J.-S. Jang (2026). A study on methods for building an official-office database from the Seungjeongwon Ilgi using named entity recognition and relation extraction. *Humanities Contents* (80), 55–84.

Shi, K. (2023). Buying public offices, bureaucratic diversity, and economic development. Working paper, California Institute of Technology.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics 87*(3), 355–374.

Stewman, S. and S. L. Konda (1983). Careers and organizational labor markets: Demographic models of organizational behavior. *American Journal of Sociology 88*(4), 637–685.

Strang, D. and J. N. Baron (1990). Categorical imperatives: The structure of job titles in California state agencies. *American Sociological Review 55*(4), 479–495.

Sturgis, P., T. S. Robinson, L. Fung, and C. Roberts (2026). SOCbot: Using large language models to dynamically measure and classify occupations in surveys. *Sociological Methods & Research*.

van Leeuwen, M. H. D., I. Maas, and A. Miles (2002). *HISCO: A Historical International Standard Classification of Occupations*. Leiden: Leiden University Press.

Wan, W., C. B. Ge, M. C. Friesen, S. J. Locke, D. E. Russ, I. Burstyn, C. J. O. Baker, A. Adisesh, Q. Lan, N. Rothman, A. Huss, M. van Tongeren, R. Vermeulen, and S. Peters (2023, 06). Automated coding of job descriptions from a general population study: Overview of existing tools, their application and comparison. *Annals of Work Exposures and Health 67*(5), 663–672.

White, H. C. (1970). *Chains of Opportunity: System Models of Mobility in Organizations*. Cambridge, MA: Harvard University Press.

Wu, Y. (2013). *Zhongguo de juanna zhidu yu shehui* [China's juanna system and society]. Nanjing: Jiangsu People's Publishing House. In Chinese; Japanese original, Kyoto University Academic Press, 2011.

Yan, B. (2001). "Pinwei–zhiwei" shijiao zhong de chuantong guanjiezhi wuqi yanhua [The five-stage evolution of the traditional official rank system from the perspective of rank-in-person and rank-in-position]. *Lishi yanjiu [Historical Research]* (2), 3–14. In Chinese.

Yan, B. (2010). *Zhongguo gudai guanjie zhidu yinlun* [An introduction to the official rank system of imperial China]. Beijing: Peking University Press. In Chinese.

Yang, A., A. Li, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Gao, C. Huang, C. Lv, C. Zheng, D. Liu, F. Zhou, F. Huang, F. Hu, H. Ge, H. Wei, H. Lin, J. Tang, J. Yang, J. Tu, J. Zhang, J. Yang, J. Yang, J. Zhou, J. Zhou, J. Lin, K. Dang, K. Bao, K. Yang, L. Yu, L. Deng, M. Li, M. Xue, M. Li, P. Zhang, P. Wang, Q. Zhu, R. Men, R. Gao, S. Liu, S. Luo, T. Li, T. Tang, W. Yin, X. Ren, X. Wang, X. Zhang, X. Ren, Y. Fan, Y. Su, Y. Zhang, Y. Zhang, Y. Wan, Y. Liu, Z. Wang, Z. Cui, Z. Zhang, Z. Zhou, and Z. Qiu (2025). Qwen3 technical report.

Yang, C. (2022). A new estimate of Chinese male occupational structure during 1734–1898 by sector, sub-sector pattern, and region. *The Economic History Review 75*(4), 1270–1313.

Yang, G. (2021). "Tianxia you shi gongming duo tu" he wan Qing keju rushi de bize ["Many routes to office when the realm is in turmoil" and the narrowing of entry into office through the examinations in the late Qing]. *Qingshi yanjiu [Studies in Qing History]* (2), 1–18. In Chinese.

Yu, Y., Y. Hou, Y. Wu, and C. Campbell (2026). A machine learning approach for nominative record linkage in Chinese historical databases. *Historical Methods: A Journal of Quantitative and Interdisciplinary History 59*(3), 149–166.

Zhang, L. (2013). Legacy of success: Office purchase and state-elite relations in Qing China. *Harvard Journal of Asiatic Studies 73*(2), 259–297.

Zhang, L. (2022). *Power for a Price: The Purchase of Official Appointments in Qing China*. Cambridge, MA: Harvard University Asia Center.