



Getting Down to
FACTS



Assessing Local Control and Accountability Plans (LCAPs) Using Generative AI

Jacob Hibbel
University of California Davis

Xander Beberman
Stanford University

May 2026



Stanford | SCALE Initiative
Accelerator for Learning

Assessing Local Control and Accountability Plans (LCAPs) Using Generative AI

Jacob Hibel, *University of California Davis*

Xander Beberman, *Stanford University*

Introduction

In addition to simplifying California public school funding calculations, the Local Control Funding Formula (LCFF) departed from a state-mandated accountability system wherein legislators dictated public schools' expected goals and progress. Premised on Governor Jerry Brown's adherence to the principle of subsidiarity (Bae and Stosich 2018; Wright 2017), LCFF designated local educational agencies (LEAs) as the primary budgetary decision-making units that must strategically plan future expenditures before receiving annual funds from the state. With input from community stakeholders, LEAs must write and publish a Local Control and Accountability Plan (LCAP) each academic year. In their LCAPs, LEA administrators identify goals, actions, services, and predicted expenditures aligned with statewide education priorities to improve outcomes for all students, and especially higher-need groups like economically disadvantaged students, English learners, and foster youth.

Serving simultaneously as a tool for education planning, budgeting, and accountability, LCAPs are information-dense documents that are intended to publicly record LEAs' educational priorities and resource allocation decisions. However, several challenges have rendered the information contained in LCAPs minimally accessible for large-scale, systematic research.

The primary barrier to systematic retrieval and analysis of LCAP data stems from the format in which the documents are published. LCAPs originate as templates provided to LEAs by the state as fillable tables and text fields in Microsoft Word format. LEAs complete these templates with considerable variation in length and formatting, and then publish the documents as PDF files on their organization's webpage. Early attempts to programmatically extract LCAP data relied on optical

character recognition (OCR) applied to PDF files, a method that works reasonably well for clean, typed documents but struggles with the inconsistencies introduced by district-level formatting choices, multi-column layouts, embedded tables, and scanned pages. Other text extraction tools, such as R's `pdftools` package, are reasonably effective in recovering text from digitally produced (as opposed to scanned) pdf documents, but neither text extraction tools nor OCR are able encode document structure in a reliable way. As a result, analysts may be able to extract raw text from LCAPs, but are left with no indication of the relationships among the goals, actions, and expenditures that give LCAP data their greatest analytic value.

Some researchers and policy organizations have employed manual or “hand coding” approaches to LCAP data extraction, which can yield high-quality, reliable data tailored to specific research questions. However, hand coding is comparatively labor-, time-, and resource-intensive, all of which make the manual data extraction approach best suited to studies of relatively small LCAP samples as opposed to more comprehensive analyses of California’s public education landscape.

Despite LCAPs’ significance as planning tools and accountability documents, the data challenges described above have led to a scenario in which we know relatively little about what California’s LEAs actually write in their LCAPs. Do districts in different geographic contexts or with different student populations articulate different goals and priorities? Do the goals set out by charter school LEAs look different from those described by traditional public school districts? How do the actions districts describe connect to their stated goals, and to what extent do districts with similar goals vary in the sorts of actions they propose? The necessary data for answering these sorts of questions exists within publicly available LCAPs, but extracting such information in usable form has been prohibitively difficult.

The parsing tool described in the following section addresses this gap directly. By leveraging large language models to convert LCAP PDFs into structured JSON, the tool makes it possible to extract goal text, action descriptions, and associated metadata from large document collections, thereby unlocking a corpus of LCAP text suitable for statewide education policy analysis.

The LCAP Parsing Tool

We developed the following multi-stage approach to LCAP processing and text conversion:

Stage 1: LCAP schema encoding. The outline for the LCAP template provided to LEAs was converted into a YAML schema, including major and minor section headers and table schemas. This schema serves as the reference structure against which parsed document content is later matched and validated.

Stage 2: OCR and document parsing. Since many LCAPs contain a mix of page orientations, each page of the LCAP is processed with `Tesseract` OCR to detect its orientation, and then rotated to the correct orientation and resized for OCR. Each page image is processed with DeepSeek's OCR model, which returns a structured, markdown-like representation distinguishing text, headings, and tables (with tables in HTML format). We selected DeepSeek for this stage after comparing its accuracy against Google Gemini and Claude on LCAP documents, finding it more accurate and less expensive for table extraction. Pre-processing then cleans the HTML output, converts tables to a Python-native format, and concatenates consecutive tables with matching headers or shapes.

Stage 3: LLM-based cleaning and matching. OCR artifacts and idiosyncratic formatting across individual LCAPs, such as multi-page tables with inconsistent header repetition, or rows split across pages, prevent direct structuring of the parsed output. To address this, a large language model (Google Gemini¹) is applied in two targeted passes per document. The first pass matches document section headers to expected headers derived from the LCAP template schema, returning a JSON mapping. The second pass processes each table, returning the appropriate schema, any column name corrections, and instructions for removing or merging rows to fix pagination artifacts. This two-pass design was chosen to minimize the risk of hallucinated changes to table text. When tables remain unparseable after these passes, a third request asks the model to concatenate and return corrected full tables.

¹ Google released new models during the process of parsing these PDFs, and as a result these requests were made with a mix of Gemini 2.5 Pro, 3.0 Flash, and 3.0 Pro

Stage 4: Post-processing and JSON assembly. After applying model responses, post-processing separates numbered goal sections, distinguishes major and minor document sections, and concatenates text within sections. The result is assembled into a final JSON file per LCAP, and text fields are then extracted into a document-level CSV file.

We applied this parsing method to a sample of 2,680 PDF documents comprising the school district, County Office of Education (COE), and individual charter school LCAPs submitted to the California Department of Education for the 2024-2025 academic year. We then conducted a statewide descriptive analysis focused on the content of the “Goals and Actions” LCAP section, which contains LEAs' stated goals, associated actions and services, and expenditure information. This information had previously been particularly challenging to recover from the LCAP template’s series of multi-page tables and interspersed text fields. Additionally, we merged the LCAP data with administrative records from the California Department of Education and the National Center for Education Statistics to provide key contextual covariates for the analyses described below.

Analytic Approach

This study employs three complementary analytic strategies: structural topic modeling to characterize the substantive, thematic content of LEAs’ LCAP goals, a dictionary-based typology to categorize the actions associated with each goal, and a set of elaboration measures to assess the quality and measurability of the goal descriptions. Together, these approaches illustrate a few of the many possible analyses that the parsed LCAP corpus makes possible.

Structural Topic Modeling

To identify the substantive themes present in LCAP goal text, we apply structural topic modeling (STM; Roberts et al., 2019), a type of probabilistic topic modeling that estimates the prevalence of latent topics across documents while allowing document-level covariates to influence topic prevalence estimates. STM is well suited to LCAP goal text because we anticipate that LEAs’ goals vary systematically by their organizational and contextual characteristics (e.g., district size, urbanicity, charter status). STM allows these differences to be modeled directly.

We fit the district STM on a corpus of 4,670 cleaned and deduplicated goal texts (i.e., a corpus comprising substantively meaningful goals in which each unique goal text was recorded once to avoid artificially inflated influence caused by texts appearing verbatim across multiple LEAs' LCAPs, a phenomenon we examine in depth below). Prior to fitting the model, goal text was preprocessed using standard procedures: tokenization, stopword removal, and removal of low-frequency terms. We fit models across a range of topic counts (K) and selected $K=25$ based on a consideration of diagnostic statistics including held-out likelihood, exclusivity, and semantic coherence, along with our assessment of the resulting topics' interpretability upon manual review. We assigned each topic a descriptive label based on its highest-probability terms and most representative goal texts for each topic.

Each topic's prevalence was modeled as a function of four covariates: charter network membership (Non-Charter, Independent Charter, Aspire, KIPP, Alliance, Green Dot, PUC, and Other LAUSD-Authorized Charter), urbanicity (City, Suburb, Town, Rural), log-transformed number of schools, and LCAP goal type (Broad Goal, Focus Goal, Equity Multiplier Focus Goal, Maintenance of Progress Goal). Prevalence effects were estimated using the `estimateEffect` function in the `stm` R package.

We fit a separate STM on a corpus of 272 deduplicated COE goal texts. Given the smaller corpus size, we selected $K=7$ following the same model selection criteria. The same covariates were used where applicable. COE findings are presented separately given the distinct role COEs play in the LCAP system relative to school districts and individual charter schools (Warren 2016).

Dictionary-Based Action Typology

To characterize the types of actions LEAs describe in their LCAPs, we developed a dictionary-based typology applied to the full set of parsed action texts. The typology comprises 14 categories, such as Professional Development, Staffing/Hiring, Curriculum/Materials, Technology, Family/Community Engagement, Assessment/Data Use, and Student Support Services, derived iteratively from a review of action text across a stratified sample of LEAs. Each action was tagged with any matching category, allowing actions to receive multiple tags. Approximately 7.4% of district actions and 10.4% of COE actions did not match any category and are treated as untagged in the analysis. This

untagged rate, which we discuss at greater length below, reflects both genuine ambiguity in action language and the limitations of dictionary-based methods.

Elaboration Measures

To assess variation in goal quality and specificity, we compute four elaboration measures at the goal level: word count, a measurability score reflecting the presence of outcome-oriented language, an indicator for whether the goal includes a numeric target, and a flag for whether the goal text is substantively meaningful (as opposed to a minimal placeholder or “stub”). These measures allow us to examine whether goal elaboration varies systematically by charter network, locale, or goal type, and complement the topical analysis by capturing how fully LEAs describe their goals in the LCAP.

Goal-Action Connections

Finally, we examine the relationship between goal content and action types by linking STM topic assignments to the tagged action corpus. For each goal assigned a dominant topic, we aggregate the action types associated with that goal to produce a topic-by-action-type matrix. This analysis allows us to examine whether LEAs that write goals on particular themes (e.g., closing achievement gaps, increasing college and career readiness) tend to articulate systematically similar or different types of actions for achieving those goals.

Data and Corpus Description

This study draws on a corpus of 2,680 LCAP documents submitted to the California Department of Education for the 2024–2025 academic year, spanning school districts, County Offices of Education (COEs), and individual charter schools. Of these, 2,192 documents were actual LCAPs from which relevant information could be recovered. Table 1 summarizes the composition of the corpus by LEA type.

Table 1. Corpus Composition by LEA Type

LEA Type	N LEAs	Goals (Total)	Goals (Deduplicated)
District	864	3,207	2,904
COE	58	272	272
Charter	1,270	3,883	2,365
Total	2,192	7,362	5,541

The corpus includes 864 school districts, comprising elementary, high school, and unified districts, along with 58 COEs and 1,270 individual charter schools. Together, these LEAs specified a total of 7,362 goals, of which 5,541 remained after deduplication (i.e., removing goal texts that appeared verbatim across multiple LCAPs). District goals account for 3,207 of the total (2,904 deduplicated); COE goals account for 272 (272 deduplicated); and charter goals account for 3,883 (2,365 deduplicated).

Urbanicity varies considerably across the school district sample. Rural districts are the most numerous (346), followed by Suburban (266), Town (134), and City (117) districts — a distribution that reflects California's large number of small, geographically dispersed districts. Districts also vary considerably in size, from single-school districts to large urban systems with over 200 schools. The median district operated 6 schools, though the mean of 9.5 (SD = 14.3) reflects a right-skewed distribution driven by a small number of large districts, for which we correct by including a log-transformed school count in our models.

Charter schools in the corpus include both independent charters and schools affiliated with charter management organizations (CMOs). Table 2 shows the distribution of goals and LEAs by charter network status in the STM analysis corpus. Non-charter districts account for the largest share (792 LEAs), followed by Independent Charters (715). Among the charter networks, Other LAUSD-Authorized is the largest group (151 LEAs), followed by Aspire (33), Alliance (23), KIPP (14), Green Dot (10), and PUC (9).

Table 2. Charter Network Composition (STM Analysis Corpus)

Charter Network	N LEAs	N Goals
Non-Charter	792	2,977
Independent Charter	715	2,110
Other LAUSD-Authorized	151	396
Aspire	33	99
Alliance	23	115
KIPP	14	42
Green Dot	10	49
PUC	9	54

Findings

Goal Elaboration, Quality, and Duplication

As an initial illustration of analysis enabled by the parsed LCAP corpus, we examine three dimensions of goal quality: the prevalence and patterning of verbatim goal duplication across LEAs, the length and substantiveness of goal text, and the degree to which goals employ language oriented around data and measurement. These analyses are descriptive and should be understood as proof-of-concept illustrations of the kinds of large-scale quality assessments that systematic LCAP parsing makes possible.

Goal Duplication

The LCAP corpus is notable in the prevalence of verbatim goal sharing across LEAs. Of the 6,718 goals in the district and charter school LCAP corpus, 2,470 goals (36.8%) were exact or near-exact copies of goals submitted by at least one other LEA. Duplication was heavily concentrated among charter schools, particularly those affiliated with CMO networks. As shown in Table 3, non-charter districts exhibited a goal text duplication rate of 17.8%, whereas independent charter school LCAPs showed a 48.7% goal duplication rate. Among schools belonging to named CMO networks, duplication rates were

dramatically higher: Alliance (100%), Aspire (99%), KIPP (93%), Green Dot (91.8%), and PUC (78.7%) schools filed LCAPs that were largely or entirely identical to those of their sister schools.

Table 3. Goal Duplication Rates by Charter Status and CMO Affiliation

Charter Network	N Goals	N Unique Texts	N Duplicated	% Duplicated
Alliance	115	0	115	100.0%
Aspire	99	1	98	99.0%
KIPP	43	3	40	93.0%
Green Dot	49	4	45	91.8%
PUC	61	13	48	78.7%
Other LAUSD- Authorized	450	151	299	66.4%
Independent Charter	2,506	1,286	1220	48.7%
Non-Charter	3,395	2790	605	17.8%
Total	6,718	4,248	2,470	36.8%

The most frequently repeated goal texts reflect recognizable network-level planning frameworks. Four Alliance goals, including "Scholar Culture, Staff Culture, and Engagement," "Basic Services," "College Success," and "Families as Partners," appeared across all 23 Alliance schools in the corpus, and a fifth goal, "Instructional Program," appeared across 17. Aspire's two most widely shared goals, including one centered on cultivating "inclusive, affirming, joyful, and safe learning environments" and another on "grade-level, engaging, affirming, and meaningful, culturally responsive academic programming," each appeared across 13 schools. KIPP's most prevalent goal, focused on student, family, and school engagement, appeared across 14 LCAPs, with two additional goals around student achievement appearing across 10 and 7 schools respectively.

We observed a qualitatively distinct form of goal duplication among Los Angeles County-authorized independent charters: a two-word "Academic Excellence" goal appeared across 47 separate LEAs affiliated with Independent Charter and Other LAUSD-Authorized networks. Unlike CMO goal sharing, which likely reflects centralized network-wide planning, this pattern more plausibly

reflects insufficient authorizer oversight, as these schools filed goals that were substantively empty as opposed to being substantively meaningful but uniform.

Public school districts nested within the same county share some key similarities with network-affiliated charter schools. To the extent that identical LCAP goal language at a lower organizational level arises from coordinated planning and policy-writing at a higher bureaucratic level, we might expect that goal duplication across school districts would cluster within particular COEs. To assess this possibility, we examined the subset of 216 non-charter public school district goal texts that appeared verbatim in two or more districts' LCAPs. Of these 216 goal texts, 212 (98.1%) were repeated across districts within the same county, with only four examples of cross-county goal duplication. This within-county concentration is consistent with the hypothesis that certain COEs (or county-level technical assistance providers) are supplying template goal language to multiple districts within their jurisdictions. This form of coordination would produce county-clustered duplication rather than the statewide template adoption observed among LCAPs filed by CMO network charter schools. The geographic pattern of goal duplication suggests that COE-level LCAP coordination is more prevalent in some counties than others, with Santa Clara (70 shared goal texts), Stanislaus (41), and Sonoma (34) counties demonstrating the highest concentrations of within-county goal sharing among public school districts. It remains an open question worthy of additional investigation whether this pattern reflects COEs' beneficial capacity support for smaller districts or a form of goal homogenization that runs counter to LCCFF's fundamental local control and planning intentions.

Because verbatim repeated goals dominated several STM topics in initial model fits, those models yielded topics that effectively functioned as CMO membership identifiers rather than capturing genuine variation in thematic goal content. We therefore deduplicated the goal text corpus prior to final STM fitting, retaining one copy of each unique goal description. We retained network affiliation as a prevalence covariate in the STM, allowing us to estimate how charter network membership is associated with topic prevalence across the deduplicated distribution of goal content.

Goal Length and Substantiveness

Across the deduplicated corpus, goals averaged 36.6 words (median = 28), with considerable variation ranging from single-word entries to goals exceeding 400 words. We classified goals with fewer than 10 words as stubs and excluded them from elaboration analyses; 329 goals (7.0% of the deduplicated corpus) met this criterion.

Stub rates varied substantially by charter network. Other LAUSD-Authorized charters had the highest stub rate at 12.9%, followed by Independent Charters at 8.2% and Non-Charter districts at 5.9%. Named CMO networks had stub rates of 0%, reflecting that their goals, although frequently shared across multiple LCAPs, tend to be substantively written.

Table 4. Goal Elaboration by Charter Network (Substantive Goals Only)

Charter Network	N Goals	Mean Words	Median Words	Stub Rate (<10 words)	Mean Measurability	% Numeric Target
Non-Charter	2,752	38.0	28.0	5.9%	1.11	10.1%
Independent Charter	1,357	36.5	27.0	8.2%	0.83	4.7%
Aspire	10	31.8	30.5	0.0%	1.10	0.0%
KIPP	8	77.6	81.0	0.0%	3.25	0.0%
Alliance	6	22.7	20.5	0.0%	0.17	0.0%
Green Dot	6	96.5	106.5	0.0%	5.33	50.0%
PUC	17	18.6	17.0	10.5%	0.24	23.5%
Other LAUSD-Authorized	189	39.4	29.0	12.9%	1.33	7.4%

Among substantive goals, Green Dot and KIPP stand out for goal length, with mean word counts of 96.5 and 77.6 respectively, roughly double the non-charter mean of 38.0. Alliance and PUC goals were considerably shorter, averaging 22.7 and 18.6 words respectively.

Goal Measurability

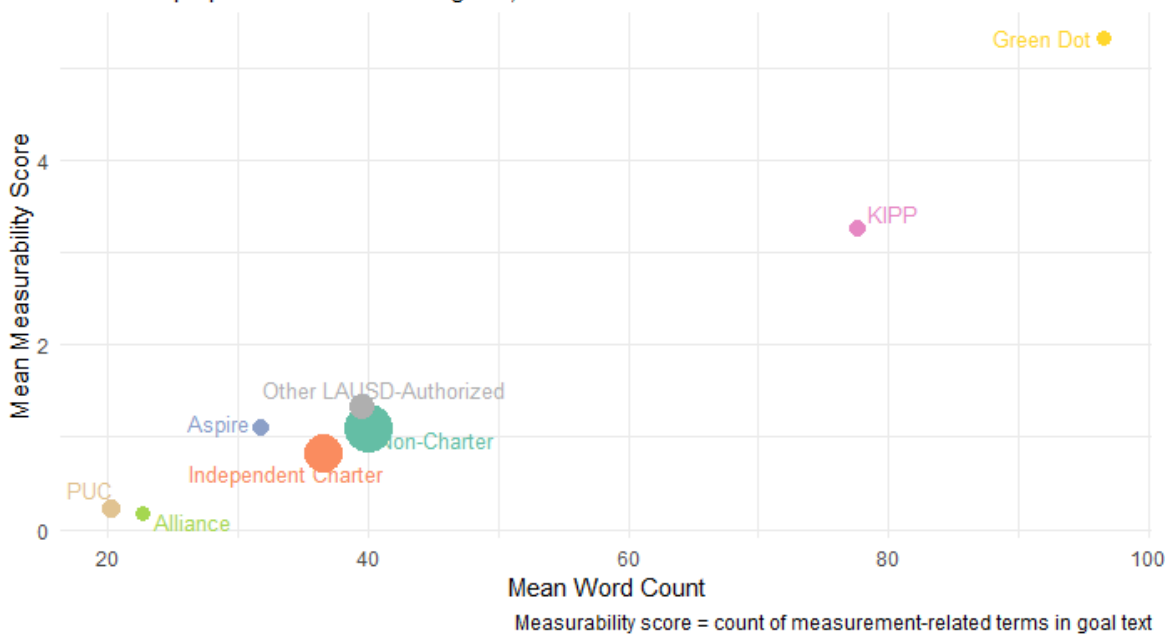
The measurability score counts occurrences of measurement-oriented terms in goal text (e.g., *assess*, *percent*, *target*, *baseline*, *proficien*) and we present it as a rough approximation of the degree to which

an LCAP’s goals emphasize quantifiable outcomes and data use. The measure should be interpreted as a fairly simple indicator of planning specificity rather than a definitive, holistic assessment of goal quality.

Overall, goals averaged a measurability score of 0.97, with a median of 0 — indicating that the majority of goals contain no measurement-oriented language at all, while a subset contain several such terms. Across the entire deduplicated corpus, only 7.9% of goals contained an explicit numeric target.

Figure 1. Goal Elaboration vs. Measurability by District Type

Point size proportional to number of goals; stubs excluded



As shown in Figure 1, LCAP goals vary substantially on both length and measurability according to charter status and CMO network affiliation. Green Dot schools’ goals are both the longest and the most measurable (mean score = 5.3, median = 7.0), with half containing a numeric target. KIPP schools’ goals are long but less measurable (mean = 3.25) and contain no numeric targets. Non-charter districts and Other LAUSD-Authorized charters cluster near the middle on both dimensions. Alliance and PUC goals tended to be short and have near-zero measurability scores, reflecting that their shared goals tend toward cultural and relational language rather than data-oriented framing.

Table 5. Goal Elaboration by Goal Type (Districts and Charters Only)

Goal Type	N Goals	Mean Words	Median Words	Mean Measurability	% Numeric Target	% Stub
Broad Goal	3,382	33.6	27.0	0.75	1.7%	6.6%
Focus Goal	433	45.2	34.0	1.76	27.5%	6.5%
Equity Multiplier Focus Goal	406	56.5	43.5	2.40	42.6%	2.5%
Maintenance of Progress Goal	339	35.4	27.0	0.58	5.3%	11.2%
Unknown	114	24.3	20.5	0.43	1.8%	27.2%

LCAP goals fit into four predetermined types: Broad Goals, Focus Goals, Maintenance of Progress Goals, and Equity Multiplier Focus Goals. Our measures of goal elaboration vary by goal type in ways that can be considered largely consistent with the intent of the LCAP framework. Focus Goals and Equity Multiplier Focus Goals are designed to concentrate attention on specific student subgroups, and their descriptions tend to be longer and more measurable than Broad Goals. Equity Multiplier Focus Goals have the highest mean word count (56.5 words), the highest measurability score (2.40), and by far the highest rate of numeric target inclusion (42.6%). Focus Goals also show elevated elaboration relative to Broad Goals on all three measures. Maintenance of Progress Goals, which document continued progress on previously met targets, have a stub rate of 11.2%, which may indicate that districts treat these as lower-priority administrative entries. Goals with unspecified type have both low elaboration and a high stub rate (27.2%), suggesting these are often incomplete or placeholder entries.

Among non-charter districts, elaboration increases modestly but consistently with district size, as shown in Table 6. Districts in the top quartile, averaging 31 schools, had a mean word count of 42.9 and a numeric target rate of 16.1%, compared to 33.0 words on average and 6.0% numeric target rate among first quartile districts, which averaged 1.3 schools. Measurability scores follow the same gradient. This pattern is consistent with differences in planning capacity, as larger districts typically have

more dedicated staff resources for LCAP development, though the relationship is not particularly dramatic.

With respect to urbanicity, city districts write the longest and most measurable goals, with a mean word count of 44.1 (median = 30.5) and a mean measurability score of 1.45, compared to rural districts at 34.5 words and 0.81 measurability. Suburban districts are comparable to city districts on word count but somewhat lower on measurability. Town and rural districts cluster together at the lower end of both dimensions. Stub rates are similar across locale categories, ranging from 4.7% (rural) to 6.9% (town), suggesting that the locale differences reflect elaboration depth rather than any systematic lack of willingness to file meaningful LCAP goals.

Table 6. Goal Elaboration by Locale and District Size (Non-Charter Districts Only)

	N Goals	Mean Words	Mean Measurability	% Numeric Target
Locale Type				
City	464	44.1	1.45	14.4%
Suburb	999	41.0	1.16	10.4%
Town	520	33.2	0.89	9.8%
Rural	943	34.5	0.81	6.0%
District Size (Quartile)				
Q1 (mean = 1.3 schools)	732	33.0	0.68	6.0%
Q2 (mean = 4.7 schools)	732	36.5	0.98	7.2%
Q3 (mean = 9.8 schools)	731	39.6	1.16	8.8%
Q4 (mean = 30.7 schools)	731	42.9	1.36	16.1%

Taken together, these patterns suggest that goal elaboration quality varies substantially across the LCAP ecosystem and that this variation is structured by organizational type, staffing capacity, regulatory intent, and, for charter schools, network affiliation. Future analyses could examine whether elaboration measures predict other outcomes of interest, such as performance on state dashboard indicators or patterns of resource allocation.

Structural Topic Modeling: Districts

To characterize the substantive content of district and charter school LCAP goals statewide, we fit a 25-topic Structural Topic Model to the deduplicated and cleaned corpus of 4,670 goal texts. The model identifies latent themes in goal language and estimates how topic prevalence varies by district characteristics. As an initial demonstration of text analysis that can be facilitated by the LLM-based parsing procedure we developed, we focus here on describing LCAP goals' topic structure and highlighting selected covariate effects that illustrate the kinds of systematic variation the model can detect.

Topic Structure and Prevalence

Table 7 presents the 25 LCAP goal topics organized thematically, along with their mean prevalence (i.e., the average proportion of each goal's content attributed to the topic) and the seven highest-scoring FREX terms. FREX terms are words that are selected based on their observed frequency within a given topic combined with their exclusivity to that topic, thereby capturing the language that most distinctively characterizes each topic. Topics span the major domains of California's state priority framework, including student achievement and curriculum, equity and unduplicated pupil services, family and community engagement, school climate and culture, and operational and resource management.

The most prevalent topics reflect broad planning priorities that cut across district types: Safe & Secure Environment (7.7%), Community Partnerships (7.1%), College & Career Readiness (6.8%), Parent & Family Engagement (6.4%), and Rigorous Curriculum (6.3%). These five topics account for roughly 34% of total goal prevalence, suggesting that safety, community engagement, and academic quality are the dominant framing categories in California LCAP goal language.

Table 7. District STM Topics: Prevalence and Top FREX Terms

Topic	Label	Mean Prevalence	Top FREX Terms (Stemmed)
Student Achievement & Curriculum			
T8	College & Career Readiness	6.80%	readi, career, colleg, pathway, prepar, technic, post-secondari
T7	Rigorous Curriculum	6.30%	curriculum, rigor, quality, high-qual, instruct, core, common
T14	Basic Conditions of Learning	5.20%	materi, facil, broad, appropri, condit, basic, prioriti
T15	ELA/Math Proficiency	5.00%	art, profici, languag, mathemat, scienc, english, math
T2	Achievement Gap Closure	4.40%	gap, close, achiev, target, intervent, outcom, maxim
T3	Standards-Based Growth	3.70%	growth, demonstr, exceed, meet, toward, standard, test
T11	Grade-Level Assessment	3.60%	grade, read, monitor, level, assess, exemplari, progress
T4	MTSS & Behavioral Support	3.50%	system, multi-ti, mtss, behavior, support, util, tier
T16	21st Century Skills	3.10%	becom, world, think, compet, critic, global, citizen
Equity & Unduplicated Pupils			
T5	Equity Multiplier: Outcomes	3.50%	year, three, graduat, indic, next, collegecar, rate
T13	EL & Unduplicated Pupils	3.30%	youth, learner, disabl, homeless, ltel, incom, term
T24	Dashboard & Subgroup Targets	2.80%	score, point, distanc, ela, red, hispan, math
T6	Equity Multiplier: Funding	1.50%	multipli, fund, equiti, lcff, start, altern, independ
Family, Community & Culture			
T18	Safe & Secure Environment	7.70%	safe, environ, welcom, secur, healthi, posit, conduc
T22	Community Partnerships	7.10%	famili, communiti, partnership, partner, member, communic
T12	Parent & Family Engagement	6.40%	involv, parent, particip, activ, encourag, stakehold, input
T19	Inclusive Culture & Belonging	4.90%	feel, cultur, connect, belong, sens, inclus, thrive
T21	Student & Staff Wellness	4.20%	emot, social, physic, well-, health, whole, child
T23	Attendance & Discipline	2.90%	attend, chronic, daili, absente, expuls, dropout, reduc
Teacher Quality & Leadership			
T17	Innovation & Improvement	4.40%	excel, experi, reach, potenti, person, comprehens, enhanc
T1	Teacher Recruitment & PD	4.00%	profession, retain, recruit, effect, capac, leadership, consist
Operations & Governance			
T20	Resource & Operations Management	2.20%	resourc, oper, sustain, infrastructur, strateg, human, technolog
T10	COE/Foster Youth Coordination	1.00%	home, plan, district, coordin, project, throughout, onlin
Residual / Mixed			
T9	Mixed	1.30%	structur, voic, joy, exist, import, build, better
T25	Broad/Mixed	1.20%	work, class, improv, provid, size, counsel, increas

Note: FREX terms are selected based on their frequency within a topic combined with their exclusivity to that topic, capturing the language that most distinctively characterizes each topic relative to others.


At the lower end of the prevalence distribution, topics labeled COE/Foster Youth Coordination (1.0%), Broad/Mixed (1.2%), and Mixed (1.3%) are residual topics, capturing goals that do not fit cleanly into other thematic categories or that address specialized populations. Equity Multiplier Funding (1.5%) and Dashboard & Subgroup Targets (2.8%) are also relatively low in overall prevalence, which is consistent with these topics being concentrated in specific goal types rather than distributed broadly across the corpus.

Covariate Effects

To examine how topic prevalence varies by district characteristics, we estimated prevalence effects for charter status/network membership, urbanicity, district size, and goal type using the `estimateEffect` function in the `stm` R package. Figure 2 displays the results for charter status as a heatmap, where each cell shows the estimated difference in topic proportion between a given type of charter LEA and non-charter districts. Red cells indicate topics more prevalent in the network relative to non-charter districts; blue cells indicate topics less prevalent.

Figure 2. Charter Network vs. Non-Charter: Topic Proportion Differences

* = p < .05 | Red = higher than Non-Charter, Blue = lower



The Figure 2 heatmap reveals substantial heterogeneity in topic profiles across charter networks, with most differences concentrated in a small number of topics. The most striking pattern involves the Basic Conditions of Learning topic, which is significantly more prevalent in goals from several CMO networks relative to non-charter districts. This suggests that CMO-affiliated schools are more likely than traditional public school districts to frame LCAP goals around foundational resource provision. The Attendance & Discipline topic shows a particularly dramatic elevation for Green Dot charter schools, which may reflect a distinctive component of that network’s organizational emphasis.

Independent Charters show a modest but significant elevation on College & Career Readiness relative to non-charter districts, and significant suppression on Dashboard & Subgroup Targets and Inclusive Culture & Belonging. This pattern suggests that independent charters write more

college-focused goals but are less likely than traditional public schools to frame goals around data accountability or inclusive school culture. Other LAUSD-Authorized schools show suppression on the Dashboard topic as well, which is consistent with the elaboration findings that this group has relatively low measurability scores.

Figure 3. Topic Proportion Differences by District Size and Urbanicity

* = $p < .05$ | Red = higher than reference, Blue = lower

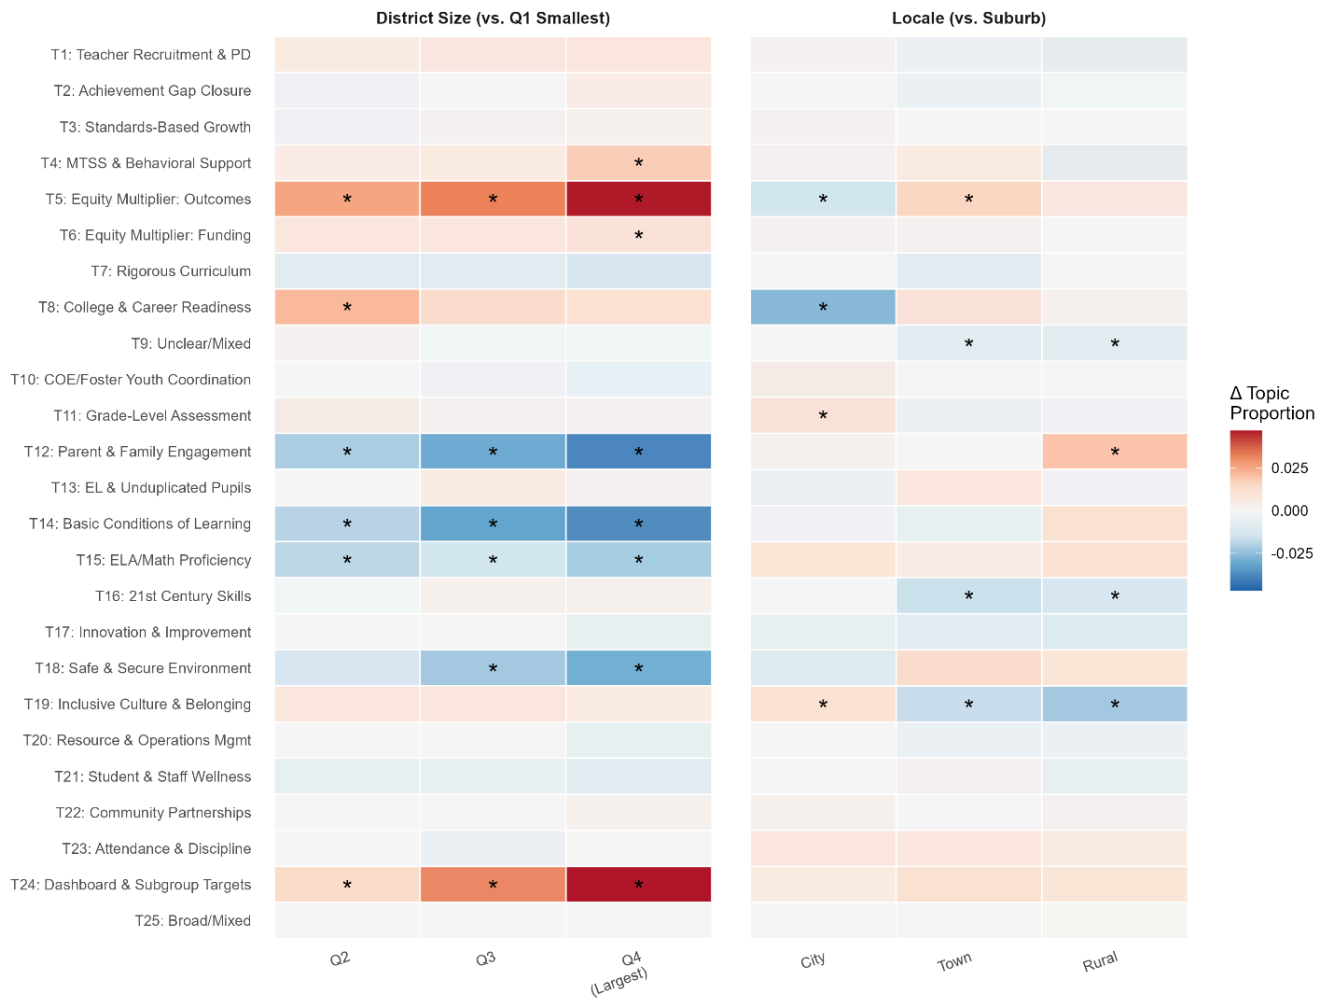


Figure 3 shows topic prevalence effects for quartiles of district size (log-transformed number of schools) and urbanicity. Larger districts show elevated prevalence for Equity Multiplier and Dashboard & Subgroup Targets, and suppressed prevalence for Basic Conditions of Learning, Parent & Family Engagement, and Safe & Secure Environment. This suggests that larger districts write more accountability- and college-focused goals while smaller districts are more likely to emphasize foundational conditions, engagement, and school safety. Urbanicity effects are more modest overall.

Inclusive Culture & Belonging is more prevalent in City districts and less prevalent in Town and Rural districts, consistent with urban districts being more likely to frame goals around belonging and cultural inclusion. Parent & Family Engagement is elevated in Rural districts, and College & Career Readiness is suppressed in City districts relative to Suburban districts, which is a somewhat counterintuitive finding.

These findings demonstrate that STM can surface meaningful and interpretable variation in LCAP goal content across district types. The patterns identified here, which include CMO networks clustering on basic conditions, larger districts emphasizing accountability, urban districts foregrounding inclusion, are broadly consistent with prior qualitative research on charter school planning and district capacity, and suggest that the parsed LCAP corpus can support systematic analysis of questions that have previously required labor-intensive hand coding.

Structural Topic Modeling: COEs

County Offices of Education occupy a distinct role in California’s LCAP system. Unlike school districts, which complete LCAPs primarily to document their own instructional programs and resource allocation, COEs serve a supervisory and support function. COEs are responsible for reviewing district LCAPs, coordinating services for foster and expelled youth, and providing technical assistance to smaller districts within their counties. To examine whether this distinct institutional role is reflected in LCAP goal content, we fit a separate 7-topic STM on the deduplicated corpus of 272 COE goals.

Table 8. County Office of Education STM Topics: Prevalence and Top FREX Terms

Topic	Label	Mean Prev.	Top FREX Terms
T1	Accountability & Data Review	3.06%	review, lea-wid, cci, perform, aco, site, input
T2	Expelled & Foster Youth Coordination	19.62%	coordin, expel, youth, foster, countywid, district
T3	Academic Achievement (ELA/Math)	15.12%	math, profici, toward, skill, intervent, ela, academi
T4	Attendance, Suspension & Subgroup Outcomes	15.05%	year, disabl, next, incom, disadvantag, hispan, socioeconom
T5	SEL, Safe Environments & Family Engagement	24.95%	famili, engag, posit, climat, health, promot, social-emot
T6	College & Career Readiness	12.50%	readi, career, colleg, cte, prepar, technic, enrol
T7	Conditions of Learning	9.71%	decreas, facil, appropri, support, standards-align, environ

Note: FREX terms are selected based on their frequency within a topic combined with their exclusivity to that topic, capturing the language that most distinctively characterizes each topic relative to others.

Topic Structure and Prevalence

Table 8 presents the 7 COE goal topics, their mean prevalence, and top FREX terms. The topic structure reflects COEs’ distinctive mandate in several ways. Expelled & Foster Youth Coordination (19.62%), which is characterized by stemmed terms like *expel*, *foster*, *youth*, *coordin*, and *countywide*, is the second most prevalent topic and has no close analog in the district model, reflecting COEs’ statutory responsibility for providing educational services to foster youth and students expelled from district schools. Accountability & Data Review (3.06%), characterized by terms like *review*, *lea-wid*, *aco*, and *cci*,

captures COEs' oversight role in reviewing and approving district LCAPs, another topic with no district parallel.


The most prevalent topic, SEL, Safe Environments & Family Engagement (24.95%), mirrors the Safe & Secure Environment topic in the district model, suggesting that school climate and safety goals are a priority across both LEA types. Similarly, Academic Achievement (15.12%), Attendance, Suspension & Subgroup Outcomes (15.05%), and College & Career Readiness (12.5%) have clear district analogs, indicating that COEs share many of the same broad planning priorities as districts even though their organizational roles differ. Conditions of Learning (9.71%) captures goals oriented around ensuring students have access to credentialed teachers, standards-aligned instructional materials, and facilities maintained in good repair, elements that map onto the conditions of learning priorities set in California's LCFF framework.

Covariate Effects

The goal type effects presented in Figure 4 are a dominant source of systematic variation in COE topic prevalence, consistent with the district findings. Attendance, Suspension & Subgroup Outcomes is strongly and significantly elevated for Equity Multiplier Focus Goals and moderately elevated for Focus Goals, patterns that mirror the district finding that equity-focused goal types concentrate on measurable outcome and subgroup accountability language. Expelled & Foster Youth Coordination is dramatically elevated for Maintenance of Progress Goals, suggesting that COEs frequently treat their expelled and foster youth coordination responsibilities as ongoing maintenance commitments rather than new or developing priorities. By contrast, Expelled & Foster Youth Coordination topics are significantly suppressed among Equity Multiplier Focus Goals, indicating that these mandated coordination functions are rarely framed as components of equity multiplier funding objectives. SEL, Safe Environments & Family Engagement is significantly suppressed for both Equity Multiplier Focus Goals and Maintenance of Progress Goals, indicating that it functions primarily as a broad institutional goal rather than a targeted or equity-focused one.

Figure 4. COE Topic Proportions by Goal Type

* = $p < .05$ | Red = higher than Broad Goal, Blue = lower

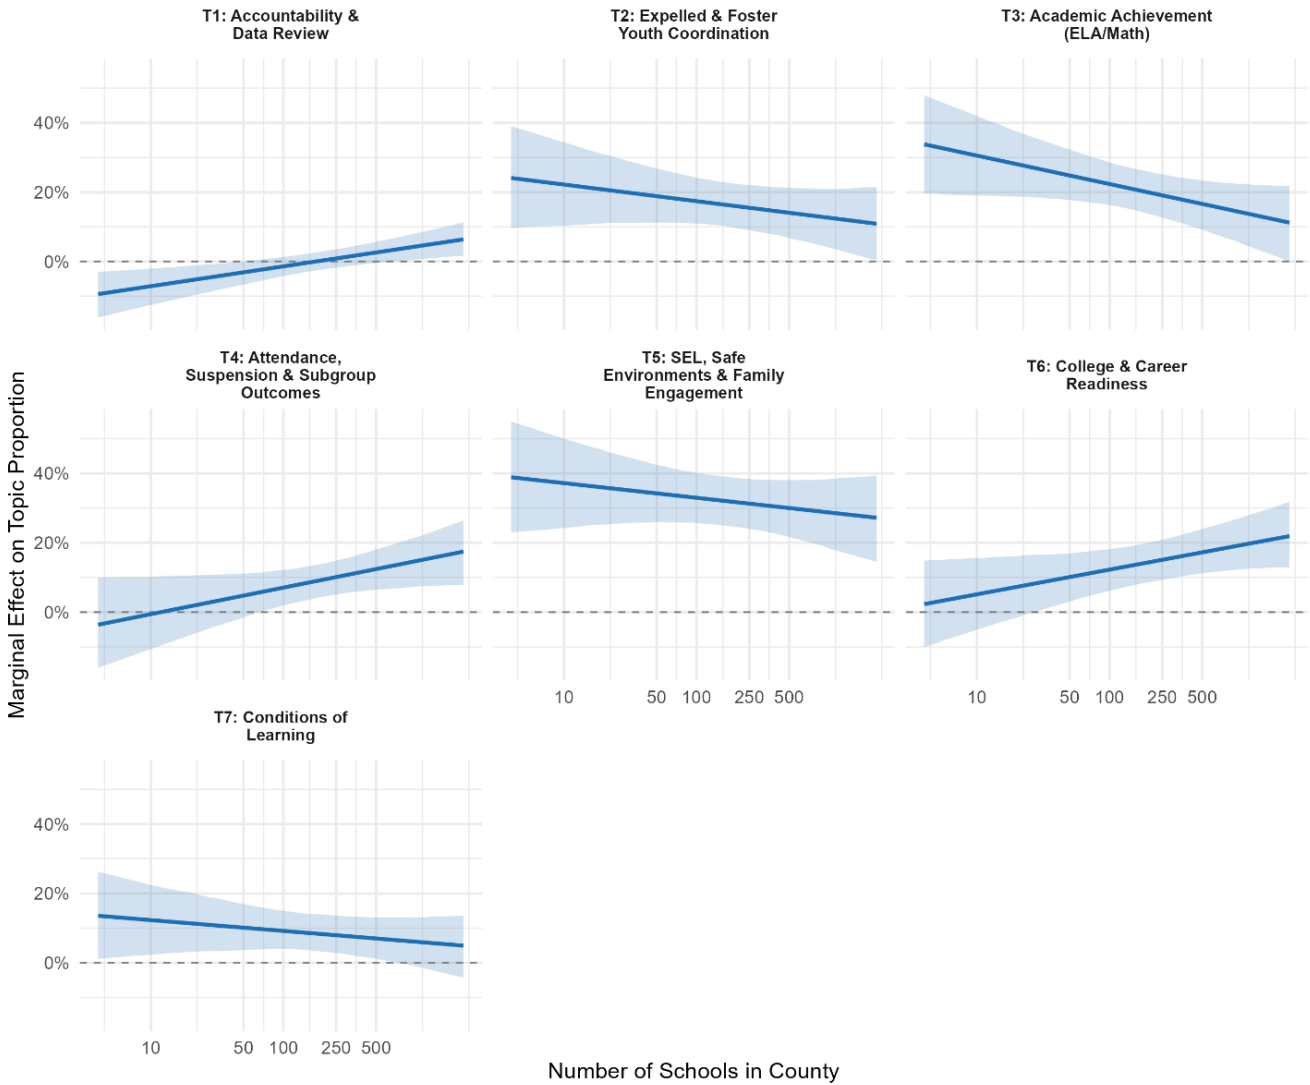


As Figure 5 illustrates, associations between county size and topic prevalence are notable for several topics. Expelled & Foster Youth Coordination and Academic Achievement (ELA/Math) both decrease in prevalence with the number of schools in the county, suggesting that smaller COEs are more likely to write goals oriented around their mandated coordination responsibilities and direct academic achievement targets than COEs serving more schools. Conditions of Learning’s prevalence is also negatively associated with county size, consistent with smaller COEs placing greater emphasis on foundational instructional infrastructure. College & Career Readiness and Accountability & Data Review demonstrate similar modest, positive associations with county size, which is consistent with larger COEs

writing more process- and pathway-oriented goals than smaller county offices.

Figure 5. COE Topic Prevalence by County Size


Marginal effects of log(N schools); reference = Broad Goal type, City locale



These size gradients are visible in Figure 6, which displays mean topic proportions by quartile of county size. The largest COEs (Q4) show notably higher prevalence of Accountability & Data Review and Attendance, Suspension & Subgroup Outcomes relative to smaller COEs, consistent with larger systems devoting elevated attention to formal accountability processes and subgroup outcome monitoring. Expelled & Foster Youth Coordination is most prevalent among quartile 1 and 2 COEs, suggesting that this mandated coordination function looms larger in the goal-setting of counties with fewer schools, perhaps reflecting the relatively greater salience of individual cases in smaller systems. SEL, Safe

Environments & Family Engagement remains the dominant topic across all county size quartiles, indicating that this broad institutional priority is shared by counties throughout the state.

Figure 6. COE LCAP Topic Proportions by County Size



Comparison with District Model

Taken together, the COE topic structure both mirrors and diverges from the district model in theoretically coherent ways. Topics related to school climate, family engagement, academic achievement, and college readiness appear in both models, reflecting shared planning priorities across LEA types. But the COE model surfaces topics associated with foster youth coordination, accountability monitoring, and county administrative relationships that are specific to the COE role and largely absent from district goal language. This divergence validates our analytic approach of modeling COEs separately, as a combined model would likely obscure the distinctive institutional content of COE goals within the much larger district corpus.

Action Typology

Whereas we used an STM analysis to characterize the content of LCAP goals, a closely connected question concerns the types of actions districts plan to take in pursuit of those goals. To investigate these LCAP actions, we applied a dictionary-based typology to the full corpus of LEA actions, classifying each action into one or more of 14 categories based on keyword matching against the action description text. Because actions can match multiple categories, the categories are not mutually exclusive and shares sum to more than 100% across the full corpus. Approximately 7.4% of district actions were “untagged,” meaning they did not match to any category, reflecting the limits of keyword matching for actions written in highly idiosyncratic or abbreviated language.

Action Category Definitions and Prevalence

Table 9 defines the 14 action categories and presents their overall prevalence in the district/charter LEA corpus. Enrichment & Extracurricular stands apart as the most common category, appearing in 42% of all tagged district actions. Staffing & Hiring (34%), Student Support Services (31%), Family & Community Engagement (27%), and Professional Development (26%) round out the top five action categories, consistent with the goal-level finding that these themes dominate California LCAP planning. English Learner supports (21%) are also a prevalent type of action, reflecting California's strong policy emphasis on EL achievement. At the lower end, Technology (11%) and Attendance & Engagement (11%) are the least common action types, though their low overall prevalence may reflect that these activities are more often described as components of broader actions rather than standalone commitments.

Table 9. Action Typology: Category Definitions and District Prevalence

Category	Keywords / Indicators	District Prevalence
Enrichment & Extracurricular	Arts, music, PE, after-school programs, CTE, TK, early childhood, ELOP, STEAM	42%
Staffing & Hiring	Hiring, recruitment, FTE, personnel, classified staff, counselor, assistant principal	34%
Student Support Services	Counseling, mental health, SEL, MTSS, RTI, intervention, tutoring, restorative practices	31%
Family & Community Engagement	Parent involvement, community events, outreach, Title I meetings, family nights	27%
Professional Development	PD sessions, coaching, PLCs, instructional coaching, capacity building	26%
English Learner	ELD, EL supports, multilingual, RFEP, ELPI, English proficiency	21%
Curriculum & Materials	Curriculum adoption, textbooks, supplemental materials, core curriculum	21%
Assessment & Data Use	Benchmark assessments, CAASPP/SBAC, iReady, DIBELS, data cycles, progress monitoring	21%
Subject Instruction	Math, ELA, literacy, reading, writing, science, history, phonics, algebra	20%
Facilities & Operations	Maintenance, repair, campus infrastructure, construction, modernization, utilities	15%
College & Career	College pathways, AVID, dual enrollment, A-G, graduation, CTE, STEM, post-secondary	13%
Special Education	IEP, SPED, RSP, paraprofessionals, inclusion, 504, disabilities	12%
Attendance & Engagement	Chronic absenteeism, SART, transportation, bus pass, tardy	11%
Technology	Devices, Chromebooks, software, digital platforms, hotspots, bandwidth	11%

As with goal descriptions, we analyzed the COE LCAPs' action typology separately, owing to county offices' distinct roles and responsibilities. The COE action distribution closely mirrors the district pattern with one notable difference: Staffing & Hiring is the most prevalent COE action type (33.4%), edging out Enrichment & Extracurricular (30.3%), whereas among districts Enrichment &

Extracurricular is the most prevalent type (42.1%) followed by Staffing & Hiring (34.1%). Additionally, College & Career actions are substantially more prevalent among COEs (19.7%) than districts (13.2%), which is consistent with COEs' role in supporting college and career transition programming and coordinating career and technical education (CTE) services across their counties.

Action Types by Goal Topic

Our final set of analyses concerns the links between districts' and charter schools' LCAP goals and their associated actions. Figure 7 presents a heatmap of action type prevalence by dominant goal topic, where each cell shows the share of actions associated with a given topic that match each action type.

Several topic-action relationships are particularly notable. Enrichment & Extracurricular actions are clearly the most popular action type, regardless of goal topic. This action type, which is indicated by words such as *arts, music, PE, after-school programs, CTE, TK, early childhood, ELOP, and STEAM*, seems to offer LEAs a flexible and broadly applicable toolkit for achieving a diverse range of objectives. By contrast, and perhaps surprisingly, Technology actions (characterized by terms such as *devices, Chromebooks, software, digital platforms, hotspots, and bandwidth*) are not prevalent in any of the 25 goal topics.

The Community Partnerships goal topic shows the highest concentration of any topic-action pairing in the matrix, with 70% of its associated actions tagged as Family & Community Engagement. The Parent & Family Engagement goal is similarly concentrated, with 57% of its actions in the same Family & Community Engagement category. The Teacher Recruitment & PD goal is strongly linked to Staffing & Hiring (46%) and Professional Development (43%) actions, and Student & Staff Wellness goals show a corresponding link to actions tagged as Student Support Services (48%) and Staffing & Hiring (41%). These highly coherent topic-action pairings suggest that, for certain goal types, districts are translating goal content into action investments in a direct and consistent way.

However, other goal topics show more diffuse action profiles. Innovation & Improvement, Rigorous Curriculum, and Basic Conditions of Learning goals each show moderate shares across many action types with no single category dominating, which may stem from their role as broad programmatic goals toward which districts apply a diverse range of implementation strategies. Safe &

Secure Environment goals are backed primarily by Staffing & Hiring (37%), Student Support Services (35%), and Family & Community Engagement (33%) actions rather than by Facilities & Operations (23%), suggesting that safety goals in California LCAPs are more frequently pursued via personnel and community investments than through physical campus investments.

Figure 7. Action type prevalence by dominant goal topic



Two topic-action patterns are particularly notable from a policy standpoint. Equity Multiplier: Outcomes is the only LCAP goal topic in which College & Career actions rank first (38%), ahead of Enrichment & Extracurricular actions, suggesting that districts framing equity goals around outcomes are making concerted investments in college pathway supports. By contrast, the College & Career Readiness goal topic has English Learner actions at only 28%, suggesting that college readiness goals are not systematically paired with EL-specific implementation strategies, even in a state in which EL students represent a large share of the college-going population.

The Dashboard & Subgroup Targets goal topic has the highest share of Assessment & Data Use actions of any topic (24%), consistent with its accountability framing, yet Enrichment & Extracurricular (45%) and Student Support Services (40%) remain its top two action types. Thus, even the most accountability-oriented LCAP goals tend to be supported primarily by programmatic investments rather than actions designed to improve monitoring and measurement.

Table 10. Mean and Median Actions per Goal by Dominant STM Topic

Topic	Mean Actions	Median Actions	N Goals
T2: Achievement Gap Closure	8.6	7	172
T7: Rigorous Curriculum	8.4	7	268
T16: 21st Century Skills	8.4	7	146
T8: College & Career Readiness	8.1	6	402
T17: Innovation & Improvement	8.0	6	130
T19: Inclusive Culture & Belonging	7.7	7	230
T3: Standards-Based Growth	7.7	6	137
T13: EL & Unduplicated Pupils	7.5	7	161
T21: Student & Staff Wellness	7.4	6	200
T4: MTSS & Behavioral Support	7.4	6	135
T15: ELA/Math Proficiency	7.3	6	260
T20: Resource & Operations Mgmt	7.2	5	75
T11: Grade-Level Assessment	7.0	6	155
T18: Safe & Secure Environment	6.9	6	431
T14: Basic Conditions of Learning	6.6	5	270
T1: Teacher Recruitment & PD	5.6	4	161
T12: Parent & Family Engagement	5.4	5	312
T22: Community Partnerships	5.2	4	347
T23: Attendance & Discipline	5.1	4	146
T6: Equity Multiplier: Funding	4.9	3	69
T10: COE/Foster Youth Coordination	4.1	3	26
T24: Dashboard & Subgroup Targets	4.1	4	143
T5: Equity Multiplier: Outcomes	4.0	3	172

Note: Topics are from the district/charter LEA STM (K=25). T9 Mixed (n = 24) and T25 Broad/Mixed (n = 2) excluded due to small sample sizes. Actions can match multiple type categories; counts reflect total actions per goal regardless of type.

Action Counts by Goal Topic

Beyond the types of actions to which districts commit, the number of actions per goal provides an indication of how extensively a goal is resourced. Table 10 presents mean and median action counts by dominant STM topic. The variation is substantial: goals characterized by the Achievement Gap Closure topic average 8.6 actions each, while Equity Multiplier: Outcomes and Dashboard & Subgroup Targets goals average only 4.0 and 4.1 associated actions, respectively.

This pattern has a notable policy implication. Two goal topics that are explicitly oriented toward measurable outcomes and accountability, Dashboard & Subgroup Targets and Equity Multiplier: Outcomes, are among the least supported as measured by action count, while topics more closely associated with instructional improvement and student supports like Achievement Gap Closure, Rigorous Curriculum, and College & Career Readiness average more than eight actions per goal. This divergence may indicate that accountability-focused goals in California LCAPs function more as monitoring commitments than as programmatic interventions. Teacher Recruitment & PD (5.6 actions) and Parent & Family Engagement (5.4 actions) also fall below the overall corpus mean, suggesting that although engagement and staffing goals are quite prevalent across the LCAP landscape, they tend to be supported by a narrower set of implementation actions than instructional improvement goals.

Viewed together, these analyses illustrate that the parsed LCAP corpus is well suited to support linkages between goals and actions, allowing researchers to assess what LEAs state they intend to achieve alongside the steps they intend to take to achieve those goals. This is precisely the kind of LCAP data linkage that has not been previously possible at a large scale.

Discussion

We prepared this technical report with two interrelated aims: to introduce a scalable LLM-based pipeline for parsing California LCAPs into structured, analyzable data, and to demonstrate the empirical potential of that corpus through a descriptive study of LCAP goal content, action typology, and goal-action linkages. The findings provide an initial illustration of what is possible when LCAP data are made computationally accessible, and offer a sense of what those data reveal about how California

districts, charter schools, and COEs translate state planning requirements into explicitly articulated commitments.

Several patterns in the data may prove interesting to policymakers, practitioners, and researchers. First, the pervasive duplication of LCAP goal text among charter management organization schools raises questions about the degree to which the LCAP process is serving its intended purpose as a genuine, "locally controlled" planning tool among network-affiliated schools, or whether it could more accurately be described as a compliance ritual in which these schools' templates are populated with externally supplied goal and action language. The discovery of 47 identical, minimally described "Academic Excellence" goals across LA County-authorized charters points to a potential oversight gap: LCAP authorizers may not be scrutinizing goal content with sufficient rigor to detect this kind of wholesale template adoption.

The within-county clustering of goal duplication among non-charter school districts raises a parallel set of questions about the role of COEs in district LCAP development. The finding that 98.1% of duplicated goal texts are shared exclusively among districts within the same county is consistent with COEs supplying template language to their affiliated districts. This form of capacity support may be particularly valuable for small, resource-constrained districts that struggle to keep up with reporting and planning demands in addition to their other responsibilities. By the same token, within-county clustering of duplicated LCAP goals raises questions about whether the resulting LCAPs reflect genuine local, school district-level planning or county-wide homogenization. The LCAP framework was designed with the intent of giving local communities a meaningful voice in educational planning; if some COEs routinely author or "contract out" their districts' plans, the "local" in local control may in practice refer to the county level rather than the district or community level. Whether COE involvement in district LCAP development is beneficial, neutral, or problematic likely depends on the quality of the resulting goals and the degree to which community stakeholders are genuinely engaged in the process. These are the kind of questions that a parsed LCAP corpus can help future researchers address.

Second, our large-scale analysis of LCAP goal elaboration suggests that there may be additional sources of LCAP goal quality problems beyond goal duplication. With a median measurability score of zero across the full corpus and only 7.9% of goals including numeric targets, the typical California LCAP

goal provides scant basis for accountability. The Equity Multiplier Focus Goal type is a stark exception, with a 42.6% numeric target rate and the highest mean measurability scores. Perhaps the structural incentives embedded in LCFF's unduplicated pupil weighting induce more rigorous goal writing for the students the policy is designed to serve. Whether more rigorous goals (as measured by the metrics introduced in this report) translate into better student outcomes is a question the parsed LCAP corpus can now help researchers address.

Third, the topic-action analysis reveals a tension between goal framing and resource allocation that has implications for how policymakers interpret LCAP implementation. Goals oriented toward dashboard accountability and subgroup targets are supported by the fewest actions in the corpus, averaging only 4.1 actions per goal, whereas more action-intensive topics like Achievement Gap Closure and College & Career Readiness average more than eight. If accountability-focused goals are consistently backed by fewer implementation actions than instructional improvement goals, this pattern may indicate that LEAs tend to regard state dashboard requirements as reporting obligations rather than as high-priority areas for resource allocation.

Limitations

Several limitations of the current study should be acknowledged. Although the LLM-based parsing pipeline represents a substantial improvement over prior approaches to LCAP data extraction, it is not error-free. OCR failures, inconsistent PDF formatting, and variation in how districts structure their LCAPs all introduce noise into the extracted data. The 390 district actions with missing text due to parsing failures (0.8% of the corpus) are an example, but unknown errors may be distributed non-randomly across LEAs. An examination of parsing failure rates across LEA types revealed no evidence of dramatic bias: failure rates ranged from 0.45% among elementary school districts to 1.05% among charter schools, suggesting that parsing quality is high and broadly comparable across the LEA types compared in this study, though charter school actions were approximately twice as likely to have missing text as non-charter district actions (1.05% vs. 0.47%). Researchers using the LCAP corpus should assess parsing quality for their specific analytic sample before drawing conclusions, and it is

essential to treat parsing artifacts as a potential source of measurement bias rather than assuming clean extraction or random measurement error.

Additionally, the STM analyses presented here reflect a single model specification chosen on the basis of interpretability and fit diagnostics, but all topic models are sensitive to the number of topics specified and to preprocessing decisions that modify the raw text (Weston et al. 2023). The K=25 district/charter model and K=7 COE model produced interpretable and substantively coherent topic structures that were supported by favorable held-out likelihood, semantic coherence, and exclusivity diagnostics, but alternative specifications would inevitably yield different topic configurations and potentially different covariate effect patterns. Therefore, these findings should be understood as one reasonable characterization of the latent thematic structure of LCAP goal language rather than as a definitively “correct” topical decomposition. Replication with alternative K values and preprocessing choices would strengthen confidence in the patterns reported here. Similarly, other topic modeling specifications, including Latent Dirichlet Allocation (LDA) and correlated topic models, could yield different topic structures and would serve as useful robustness checks.

Like topic modeling, the dictionary-based action typology has known coverage limitations. Approximately 7.4% of district/charter LCAP actions and 10.4% of COE actions did not match any of the 14 action type categories, and actions can match multiple categories simultaneously. The untagged actions are not a random sample; they likely include idiosyncratic local initiatives that fall outside the dictionary's vocabulary as well as abbreviated or poorly written actions that are not amenable to keyword matching. Analyses based on this (and, indeed, any) rule-based typology should carefully consider the possibility that systematic variation in tagging rates across LEA types could produce biased inferences.

Finally, all findings in this report are descriptive, and none should be interpreted causally. The STM covariate effects describe associations between district characteristics and goal content, but they by no means establish that charter network membership, urbanicity, or district size *causes* districts to write goals with particular topical profiles. In a similar vein, the action count and action type findings describe what districts say they will do in their annual reporting, and do not capture what they actually implemented, how faithfully they adhered to their stated goals and actions, nor what effect those

commitments had on students. Future work linking longitudinal LCAP content to student outcome data and detailed expenditure records will be necessary to draw any solidly evaluative conclusions.

Implications for Future Research

The parsed LCAP corpus opens several lines of inquiry that have been difficult to pursue with existing data. If applied to multiple years of LCAP data, this parsing approach would facilitate longitudinal analysis of LCAP goal content that tracks when, where, and how LEA planning changes, and to what ends. For example, longitudinal parsed data will make it possible to examine whether districts that receive poor dashboard ratings subsequently shift toward more accountability-oriented goal topics, or whether goal and action language is largely stable regardless of performance.

The complete corpus also enables comparative analysis at a scale at which hand-coding is untenable. Researchers could examine whether districts with similar demographic profiles but different outcome trajectories differ systematically in goal content or action investment, using the kind of STM topic proportions and action type distributions developed here as outcome or predictor variables in quasi-experimental designs. The goal elaboration measures developed in this study (word count, measurability score, numeric target rate) could serve as proxies for planning quality in studies examining whether LCAP rigor predicts student outcomes.

Lastly, the COE corpus opens an underexplored area of LCAP research. Although there is burgeoning interest in the role of COEs under LCFF, they have thus far received less research attention than might be expected given their statutory role in reviewing district LCAPs and coordinating services countywide. The finding that COE goal content is meaningfully distinct from district content suggests that future research on COE planning may yield new, important understandings of the California public education ecosystem.

Conclusion

LCAPs represent one of the most ambitious attempts in American education policy to require districts to publicly articulate their goals, strategies, and resource commitments in a single document. However, the structural complexity of LCAP PDF documents has made systematic statewide analysis of those plans difficult, if not impossible. The LLM-based parsing pipeline introduced in this report removes that barrier, allowing researchers to convert archives of complex and unwieldy PDF documents into a structured corpus that allows LCAP data to be incorporated into large-scale quantitative analysis.

The initial proof-of-concept analyses presented in this report demonstrate that the parsed corpus can support substantive research across multiple dimensions of LCAP quality and content. The findings are preliminary and should be interpreted with appropriate caution, but they illustrate the range of questions that a parsed LCAP corpus can address and establish a methodological foundation for future work.

We hope that a user-friendly, publicly accessible LCAP processing tool will help future California education scholarship expand beyond the case study and small-sample designs that have characterized the field to date, and enable the production and analysis of statewide evidence that California's accountability framework was designed to support. Moreover, we hope that making LCAP data more accessible will prove useful to the district administrators, COE staff, and charter school leaders who invest considerable time and effort in the LCAP process, giving them a clear, evidence-based understanding of how their own planning practices compare to those of their peers across the state.

References

- Bae, S., & Stosich, E. L. (2018). Redesigning state policy for meaningful and equitable learning: Lessons from California, Iowa, New Hampshire, and Vermont. Palo Alto, CA: Stanford Center for Opportunity Policy in Education. <https://edpolicy.stanford.edu/sites/default/files/Redesigning%20State%20Policy%20v2.pdf>.
- Santibañez, L., & Umansky, I. (2018). English learners: Charting their experiences and mapping their futures in California schools. *Policy Analysis for California Education*. https://edpolicyinca.org/sites/default/files/2024-10/gdtfii_brief_englishlearners.pdf
- Roberts, M. E., Stewart, B. M., & Tingley, D. (2019). Stm: An R package for structural topic models. *Journal of statistical software*, 91, 1-40.
- Warren, P. (2016). Strengthening Local K-12 Accountability: The Role of County Offices of Education. Public Policy Institute of California. https://www.ppic.org/wp-content/uploads/content/pubs/report/R_816PWR.pdf
- Weston S.J., Shryock I., Light R., Fisher P.A. (2023). Selecting the number and labels of topics in topic modeling: A tutorial. *Advances in Methods and Practices in Psychological Science*. 6(2). doi:[10.1177/25152459231160105](https://doi.org/10.1177/25152459231160105)
- Wright, P. (2017). Democratically governing school districts: How four California school districts responded to the Local Control Funding Formula. PhD Dissertation, Department of Political Science, University of California, Santa Barbara.