

Measuring Social Integration Through Participation: Categorizing Organizations and Leisure Activities in the Displaced Karelians Interview Archive using LLMs

Joonatan Laato¹, Veera Schroderus², Jenna Kanerva¹,
Jenni Kauppi², Virpi Lummaa², and Filip Ginter^{1,3}

¹TurkuNLP, University of Turku, Finland

²Department of Biology, University of Turku, Finland

³ELLIS Institute Finland

{joonatan.m.laato, veera.i.schroderus,
jmnybl, jejakau, virpi.lummaa, figint}@utu.fi

Abstract

Digitized historical archives make it possible to study everyday social life on a large scale, but the information extracted directly from text often does not directly allow one to answer the research questions posed by historians or sociologists in a quantitative manner. We address this problem in a large collection of Finnish World War II Karelian evacuee family interviews. Prior work extracted more than 350K mentions of leisure time activities and organizational memberships from these interviews, yielding 71K unique activity and organization names—far too many to analyze directly.

We develop a categorization framework that captures key aspects of participation (the kind of activity/organization, how social it typically is, how regularly it happens, and how physically demanding it is). We annotate a gold-standard set to allow for a reliable evaluation, and then test whether large language models can apply the same schema at scale. Using a simple voting approach across multiple model runs, we find that an open-weight LLM can closely match expert judgments. Finally, we apply the method to label the 350K entities, producing a structured resource for downstream studies of social integration and related outcomes.

1 Introduction

Digitizing large collections of historical records, combined with the use of large language models (LLMs), has opened new possibilities for analyzing cultural and social patterns at a large scale. LLMs make it straightforward to automate tasks such as extracting entities and other structured information from massive archival collections. In many cases, however, entity extraction alone is not sufficient; additional steps—such as grounding and grouping

entities—are often needed to support the underlying research.

In this study, we examine a unique historical corpus: *Siirtokarjalaisten tie* (*The Path of Displaced Karelians*), a large interview collection documenting World War II evacuees who were permanently relocated from the Karelia region to different parts of Finland. The archive contains testimonies from approximately 160,000 individuals, offering rare coverage of everyday social life at scale. The interviews are brief (see Figure 1), but they typically include information about organization membership and leisure activities. Organizations include a wide variety of different clubs, choirs, work unions, boards, and other similar organizations, while leisure activities primarily refer to hobbies, sports, outdoors activities, and similar. Together, these in many cases allow inferring the degree of social integration of the individuals, and, if extracted across the whole data, of the displaced population at large. Extracting this information is the primary focus of this study.

In prior work, we used LLMs to extract the organization membership and leisure activity mentions from the interviews (Laato et al., 2024). The resulting dataset contains 354,302 extracted entity mentions from 89,339 interviews. However, without further grounding and interpretation, these entities are not directly usable as a proxy for social integration. For instance, the popular “Martha” club of homemakers focusing on spreading household advice appears under multiple free-form variants (e.g., “Marthas” vs. “Martha Club”), and historically meaningful groups may be described with widely differing names across regions and dialects. Critically, mere mentions of organizations and leisure activities lack metadata that would help estimate the degree of social integration implied by a given activity.

This motivates a central step between extraction and domain analysis: normalization and categorization. For sociological questions—especially those focused on social integration and its links to health and longevity—researchers need a representation of participation that is comparable across individuals and communities. Achieving this requires mapping historical mentions to a consistent, interpretable classification framework that reflects culturally specific forms of Finnish social life.

To support this downstream analysis, we evaluate whether LLMs can assist with this structuring step and how closely they can match human expert judgment. In collaboration with sociologists (domain experts), we develop a classification framework designed to capture the dimensions of social participation most relevant to studying integration patterns, and we assess both human reliability and model performance within that framework.

We address three primary questions:

1. Is it feasible to represent the multi-dimensional nature of historical Finnish social organizations and leisure activities in a way that still yields reliable human consensus?
2. If so, to what extent can LLMs replicate domain-expert categorizations of culturally specific historical entities?
3. Can iterative and automated prompt engineering further improve model performance on this task?

2 Data

Our primary source “Siirtokarjalaisten tie” (The Path of Displaced Karelians) is a 4-volume book series documenting Finnish citizens permanently displaced from Eastern Karelia following the 1939-1944 war with the USSR. The volumes contain 89,339 family interviews conducted in 1968–1970 by approximately 300 trained interviewers, representing approximately 160,000 adults from the roughly 420,000 total refugees. All interviews were originally conducted in Finnish. The books were digitized through OCR by [Loehr et al. \(2017\)](#).

The interviews combine standardized demographic sections with free-text information about daily life, including social organization membership and hobbies. Figure 1 illustrates an example interview showing the information typically captured. Previously, we used OpenAI’s GPT-4-turbo LLM to extract all mentions of social entities from these texts ([Laato et al., 2024](#)), resulting in 192,388

Primary person: KALLE VIRTANEN

Spouse: Anna Inkeri Nieminen

Source text: Managing director, born 22. 7. -29 In Antrea. Spouse, Anna Inkeri née Nieminen, shop assist., born 21. 11. -29 Uusikirkko. Children: Jari Tapio -57 Riihimäki, Jouni Juhani -58 Riihimäki, Matti Mikael -61 Loppi. Residence, in Karelia: Antrea, Kaltove-denmaankylä 39, 42-44. Other residences: Tammela -42, Ikaalinen 44-, Janakkala, Loppi, Sajaniemi 47-54, Loppi, Pilpala 54-58, Riihimäki, Loppi kk. 60-. Kalle Virtanen lives with his family in their house. Managing director Virtanen is a member in Karelians of Loppi Reg. Assn. Tax Committee and Sajaniemi’s Deer Lads Regn. Assn. His hobbies include hunting and fishing. Mrs. is a member in Karelians of Loppi Reg. Assn. and serves as the secretary of the Loppi Kuparsaari Marttas. Her hobbies include handicrafts and gardening.

Figure 1: A single interview example with entities pertaining to the husband (yellow) and wife (blue) highlighted. Translated from Finnish, preserving typographic features of the original data. Note: Names and dates have been changed.

Top 10 Hobbies	
handicrafts (<i>käsityöt</i>)	— 31,780
fishing (<i>kalastus</i>)	— 20,222
literature (<i>kirjallisuus</i>)	— 13,128
gardening (<i>puutarhanhoito</i>)	— 10,519
reading (<i>lukeminen</i>)	— 9,324
hunting (<i>metsästys</i>)	— 6,287
skiing (<i>hiihto</i>)	— 5,912
outdoor activities (<i>ulkoilu</i>)	— 5,597
sports (<i>urheilu</i>)	— 4,997
handicrafts [partitive case] (<i>käsitöitä</i>)	— 4,708
Top 10 Organizations	
Karelian Society (<i>Karjalaseura</i>)	— 6,610
Martha Association (<i>Marttayhdistys</i>)	— 5,725
youth association (<i>nuorisoseura</i>)	— 2,545
Farmers’ Association (<i>Maamiessseura</i>)	— 2,476
Marthas (<i>Martat</i>)	— 1,967
Agricultural Women (<i>Maatalousnaiset</i>)	— 1,957
Lotta Svärd (<i>Lotta Svärd</i>)	— 1,930
sports club (<i>urheiluseura</i>)	— 1,252
church choir (<i>kirkkokuoro</i>)	— 1,230
elementary school (<i>kansakoulu</i>)	— 1,212

Table 1: Top 10 most frequently mentioned hobbies and organizations. Finnish terms in italics.

mentions of hobbies and 161,914 mentions of organizations, extracted with 88.8% F-score, close to human performance at the task. These mentions comprise a total of 71,874 unique entities (7,674 hobbies and 64,200 organizations), as common activities were mentioned by many individuals.

The top 10 most frequently mentioned entities (Table 1) already show very active lifestyle and organizational involvement of the Karelian popula-

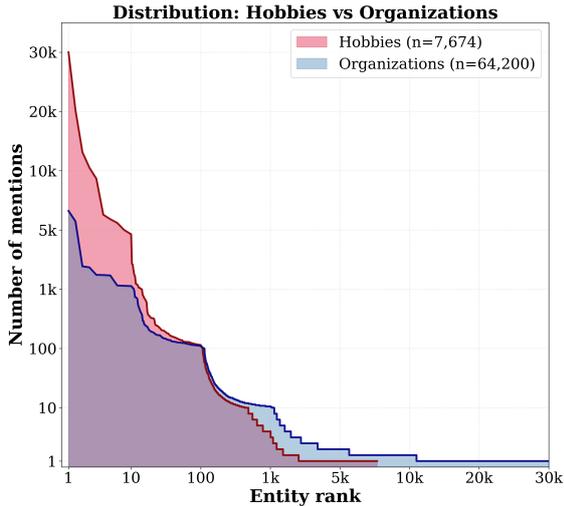


Figure 2: Distribution of hobby and organization mentions across entity ranks. Hobbies (red) show concentrated distribution in the first few mentions, with some 7.6k unique entities, while organizations (blue) show a long tail extending beyond 60k unique entities.

tion in post-war Finland. Common hobbies center on traditional rural activities such as handicrafts, fishing, and gardening. Meanwhile organizations include Karelian cultural societies, women’s associations, and agricultural groups.

As illustrated in Figure 2, the extracted hobbies and organizations display unsurprisingly different distribution. Hobbies include only 7,674 unique entities, where the top 200 account for 92.6% of all mentions (192,388 total). In contrast, organizations display a long-tail distribution with 64,200 unique entities, where the top 200 capture only 49.0% of all mentions (161,914 total).

2.1 Entity Normalization Challenges

Pattern matching reveals variations referring to effectively the same activity. The term “mart*” matches 156 unique organization strings (out of 64,200 total unique organizations) including: Martha Association (*Marttayhdistys*), Marthas (*Martat*), marthas (*martat*), Martha Society (*Marttaseura*), Martha Club (*Marttakerho*). The same root appears in 40 unique hobby strings (out of 7,674 total unique hobbies) including: Martha work (*marttatyö*), Martha activities (*marttatoiminta*). These all for the most part refer to the same popular homemaker organization and its activities.

The pattern “karjala*” (Karelia) yields 590 unique organizational strings referring to various societies established by the Karelians, often to pro-

mote their interests and maintain connections, and 94 unique hobby strings predominantly describing Karelia-specific food preparation: e.g. making Karelian pastries (*karjalanpiirakoiden valmistus*).

Handicraft variations (“käsit*”) show 78 unique hobby strings: handicrafts (*käsityöt*), handicrafts-partitive (*käsitöitä*), making handicrafts (*käsitöiden tekeminen*), handicraft making (*käsityöiden tekeminen*). The pattern also matches 22 unique organizational strings: handicraft club (*käsityökerho*).

Additionally, many organizations with entirely different names belong to the same functional category. For example, various sports clubs while having unique names should nevertheless be grouped together for the analysis.

2.2 Pre-processing: Hierarchy Extraction

We pre-process the organizations using an LLM to separate out hierarchies of organizations and individual roles: given an entity like “Chairman of the Administrative Board of the Karelian Society,” we prompt an LLM (Llama-3.3) to separate the individual’s role as well as the underlying hierarchy. In this case we would extract “Chairman” as a role and “Administrative Board” as a hierarchical element. The remaining entity is then interpreted as the actual social organization, here “Karelian Society”. This reduced the unique social organization count from 64,200 to 52,982, extracting 2,147 different roles and 709 unique hierarchies.

3 Related Work

Historical and cultural data presents unique challenges for entity normalization and clustering: (1) classifying entities into meaningful groups, (2) normalizing variant names, and (3) capturing semantic relationships through multi-dimensional categorization.

Previously, these have been approached using distributional representations such as word2vec and fastText (Mikolov et al., 2013; Bojanowski et al., 2017) to capture semantic similarities through vector clustering (Baroni et al., 2014). Transformer-based models like BERT (Devlin et al., 2019) advanced multi-label classification through contextualized representations, but require substantial domain-specific labeled data for training. More recently, LLMs have enabled clustering entities based on contextual embeddings. Huang and He (2025) transform text clustering into classification by having the LLM generate potential labels, merge

semantically similar labels, and assign entities accordingly. This label-merging approach captures underlying data structure without requiring any fine-tuning. Viswanathan et al. (2024) demonstrated few-shot clustering through LLM-generated pairwise constraints and entity canonicalization with minimal examples. Recent work on entity matching (Huang and Zhao, 2024) highlights how LLMs implement semantic equivalence beyond syntactic similarity, though challenges remain for culturally-specific entities and abbreviated names. However, multi-label classification with LLMs presents distinct challenges. Ma et al. (2025) found that LLMs suppress multiple labels during generation despite applicable categories. Lan et al. (2024) addressed this through multi-label contrastive learning. Tabatabaei et al. (2025) demonstrated successful industrial-scale deployment with custom taxonomies for hierarchical multi-label classification. Across these tasks, careful prompt engineering significantly improves LLM’s classification accuracy, with automated methods like APE (Zhou et al., 2023) and OPRO (Yang et al., 2024) emerging to generate and iteratively optimize prompts.

4 Methods

While recent work shows that LLMs can cluster semantically (Huang and He, 2025) and normalize entities (Viswanathan et al., 2024), these approaches focus on single-label classification or binary similarity judgments. However, as demonstrated in our data analysis, the diversity of historical social organizations resists discrete categorization. Attempting to single-label these entities would require hundreds of narrowly-defined categories. Instead, a multi-label approach that captures specific dimensions of each entity proves more appropriate. Furthermore, rather than forcing data into pre-existing taxonomies (as is commonly used in NER), our approach develops classification dimensions directly towards sociological research questions about social integration patterns.

4.1 A Classification Framework for Social Integration Research

The design of this framework is directly guided by the main research objective of sociologists: characterizing the social integration of Karelian refugees after relocation. Combined with further connected data on this population, such as health and mortality information, this characterization of social integra-

tion can be used to answer questions about how different levels of integration influenced mortality outcomes. Furthermore, interactions of social integration level can be measured relative to numerous factors, such as age, residential history, distance from relatives, and whether the relocation occurred individually or with the community. The dataset’s scale and demographic uniformity provide a rare opportunity to isolate the effects of post-migration social environments on individual outcomes.

While this data, and approach, will eventually be utilized broadly across different fields, initially, the main focus will be on investigating the health effects of different kinds of social integration. Consequently, key questions and category types of the annotation framework were constructed based on well-known health factors, such as physical activity (Iwasaki et al., 2001; Rhodes et al., 2017), sociality (Holt-Lunstad et al., 2010; Hansdottir et al., 2022), and mental stimulation (Yates et al., 2016; Hansdottir et al., 2022; McQuade and O’Sullivan, 2024). The frequency of activity is relevant to health in both physical activity (Iwasaki et al., 2001; Rhodes et al., 2017) and mentally stimulating leisure activities (Weziak-Białowolska et al., 2023). For non-solitary activity, frequency and group size additionally reflect the social intensity of participation.

These considerations led to the design of an annotation schema with four main aspects for each activity/organization, formulated into four questions: **Q1 - Categories:** The general type(s) of organization/activity; **Q2 – Group Size:** How many people typically participate? **Q3 – Regularity:** How often do they meet/engage? **Q4 – Physical Intensity:** What level of physical activity is involved? Each question has predefined answer options, detailed in Table 2, with special options for ambiguous cases: “Cannot be determined” applied when the judgment is impossible to make based on the information given, and “Data error” applied when the entity is not, in fact, a leisure-time activity or organization, but rather an extraction error.

To support culturally informed annotation, a questionnaire was distributed through the Karjalan Liitto ry Facebook page, where Karelian migrants and their descendants provided ratings and descriptions for approximately 60 common ambiguous entities. While these responses were not directly incorporated into the current annotation framework, they represent a valuable reference point in the refinement of the guidelines.

Q1 – Categories	Q2 – Group Size
Professional/Work	Alone
Cultural/Traditional	Small group
Religious/Spiritual	Large group
Sports/Physical activity	Non-definable
Creative/Artistic	Data error
Educational/Academic	Q3 – Regularity
Social welfare	Regular
Administrative	Occasional
Political	Continuous
General social group	Event-based
Nature-related	Non-definable
Health-related	Data error
Property/ownership	Q4 – Physical Intensity
Animal-husbandry	Intense
Cooking	Continuous
Military-related	Light
Special Hobbies	Stationary
Non-definable	Non-definable
Data error	Data error

Table 2: Questionnaire schema showing all options for each question in our final round of annotations.

4.2 Annotation Process

Four annotators, two domain experts and two machine learning specialists, annotated five iterative rounds of 50 entities, split evenly between hobbies and organizations. After each round the schema and guidelines were refined. For example, annotators should distinguish between organizations and their administrative bodies, separate “Administrative” (governance) from “Political” (party activities), and select “Cannot be determined” rather than guess when uncertain. Once a sufficient agreement had been reached, a set of 200 entities was annotated by all four annotators, and split into a development/evaluation set of 50 entities, and a final test set of 150 entities.

We measured inter-annotator agreement using pairwise Cohen’s Kappa across both evaluation (n=50) and test (n=150) sets (See Appendix Table 8 for full results). The pairwise Kappa score shows moderate-to-substantial average agreement of 0.704 (evaluation set) and 0.690 (test set). Annotator agreement is also a proxy for the difficulty of making the judgements, and the reliability of the resulting annotations. We therefore also inspect the agreement on a per-question basis. Here Q2 (group size) reached the highest agreement at 0.772/0.781 (eval/test), while Q4 (physical activity) showed the lowest, yet still moderate agreement at 0.678/0.579.

The final ground truth was established using majority vote across the four annotators (Snow et al., 2008; Sheng et al., 2008). Due to the multi-label nature of the task, a label was selected when at least

two annotators selected it (2-out-of-4 threshold). In evaluation, human annotators are scored against leave-one-out majority, excluding their own labels (3vs1), while LLM predictions will be evaluated against the full four-annotator consensus (4vs1).

In terms of F-score on the consensus annotation labels pooled across all four questions, the annotators achieved on average 78.6% F-score (test, n=150) and 80.7% F-score (evaluation, n=50) against the leave-one-out majority vote (Appendix Table 8).

4.3 LLM Annotation

With minimal editing, the annotation guidelines were crafted into the LLM prompt (see Appendix A.1 for the full prompt). We evaluated six state-of-the-art open-weight models to determine which could best replicate human performance. Additionally, using the best open-weight model, we investigated variants such as prompting all four questions simultaneously in one prompt, as compared to prompting each question individually in separate prompts, so as to establish whether the models benefit from holistic entity understanding versus focused attention on individual classification dimensions.

4.4 Models

We evaluated the following six state-of-the-art open-weight models: Qwen2.5-72B-Instruct, a 72B dense instruction model trained on 18T tokens with RLHF supporting 131k-token context (Yang et al., 2025b); Qwen3-235B-A22B, a 235B-parameter MoE (22B active) trained on 36T tokens with native chain-of-thought capabilities (Yang et al., 2025a); DeepSeek-R1-Distill-Llama-70B, a 70B student model distilled from the 671B DeepSeek-R1 teacher with RL-driven chain-of-thought (Guo et al., 2025); Meta-Llama-3.3-70B-Instruct, a December 2024 refresh with 70B parameters and 128k context (Grattafiori et al., 2024); Llama-4-Scout-17B-16E, a 17B-active/16-expert MoE enabling 10M-token contexts (Meta-AI, 2025); Mistral-Large-Instruct-2407, a 123B flagship model with 128k context (Mistral-AI, 2024). Notably, neither Qwen3 nor DeepSeek-R1, despite their chain-of-thought capabilities, outperformed Mistral-Large on this task.

The overall results of the models in terms of F-score on the evaluation set (n=50) are summarized in Table 3, demonstrating substantial differences

Model	Avg F1
Qwen2.5-72B-Instruct	65.0
Qwen3-235B-A22B	68.2
DeepSeek-R1-Distill-Llama-70B	62.5
Meta-Llama-3.3-70B-Instruct	73.0
Llama-4-Scout-17B-16E	70.4
Mistral-Large-Instruct-2407	74.6
Mistral-Large-Instruct-2407 [†]	73.6

[†]Same model, but prompted one question at a time rather than all four questions for each entity.

Table 3: Average F-scores of each model on the evaluation set (n=50).

between individual models.

4.5 Prompt Optimization

Among the six models evaluated, Mistral-Large achieved the highest performance with 74.6% average F-score (Table 3). To test whether automated prompt engineering could further improve this baseline, two optimizer models (OpenAI’s o3 and Anthropic’s Claude 4 Opus) were used to generate optimized prompt variants when being informed about the errors observed in the annotation carried out with the original prompt.

To this end, the 50-entity evaluation set was divided into 5 splits (20 items for optimization, 30 items for evaluation), with each entity appearing in exactly two optimization sets and three evaluation sets. For each split, both optimizer models produced two prompt variants based on error feedback, yielding 20 optimized prompts total. Only 7 showed improvements on the held-out data, and these gains were individually relatively minor, on average 4pp F-score (when averaging only data splits with an improvement, i.e. an overestimate of any actual gain). This suggests the iterative human annotation process had already produced instructions clear enough for the model to approximate human judgment.

5 Results

The entire framework was evaluated against the final ground truth test set of 150 entities, split evenly between hobbies and organizations. The results are summarized in Table 4, demonstrating that human annotators clearly outperform the LLM (72.7% vs. 78.6% F-score), however the difference is not even across the four annotation questions. From the 20 automatically-optimized prompts, only one out-

Method	Schema	Q1	Q2	Q3	Q4	Avg
Orig. prompt	Full	77.0	85.4	73.6	54.6	72.7
Best prompt	Full	78.2	85.8	71.7	55.9	72.9
7 prompt ens.	Coarse	80.0	87.7	76.8	62.7	76.8
7× orig. ens.	Coarse	76.9	87.4	79.9	60.0	76.1
7 prompt ens.	Full	80.0	85.0	71.4	58.9	73.7
7× orig. ens.	Full	76.9	85.9	75.3	55.0	73.3
Human	Full	82.5	88.6	73.8	69.3	78.6
Human	Coarse	82.5	90.1	81.2	73.4	81.8

Table 4: Performance of single-prompt methods, ensemble methods, and human annotators on test set (n=150) in full and simplified coarse label schema. All values are F-score.

performed the original human-engineered prompt, achieving 72.9% F-score compared to 72.7% F-score for the original, a marginal gain of 0.2 percentage points.

We further evaluated two ensemble approaches: a majority vote across the 7 optimized prompts which have previously shown an improvement on the evaluation set, and for comparison a majority vote across 7 runs of the original prompt (leveraging the fact that the LLMs sample in generation, and the output therefore is not deterministic). As seen in Table 4, the 7-prompt ensemble achieved 73.7% F-score (93.8% of human baseline), compared to 73.3% for the repeated original prompt. Simply running the same prompt multiple times showed minor improvement (73.3% vs 72.7%), suggesting that ensemble voting can partially compensate for the stochastic nature of LLM outputs. The most notable improvement came on Q1 (Categories), where the optimized ensemble achieved 80.0% F-score versus the repeated original’s 76.9%. However, the optimized ensemble performed worse on Q3 (Frequency), suggesting that different prompts may have complementary strengths and weaknesses.

5.1 Coarse Label Schema: Improving Performance Through Reduced Granularity

Considering that the downstream use of the labeled data requires a maximal attainable accuracy, we also experimented with reduced label granularity, producing a coarse schema by treating the following labels as equivalent: **Q2**: “Small group” = “Large group”; **Q3**: “Occasional” = “Event-based”; **Q4**: “Intense” = “Continuous” = “Light”. This simplification essentially reduces Q2 to the choice of “solitary vs social,” Q3 to the choice of “regular vs rare,” and Q4 to the choice “active vs stationary,”

while keeping Q1’s full categorical granularity.

With these coarse labels, the optimized ensemble achieved 76.8% F1 score, representing 93.9% of human baseline. This represents a substantial 3.1 percentage point improvement over the same ensemble with full label granularity (73.7%). The gains were most pronounced in Q2 (Group Size), where collapsing “Small group” and “Large group” into a single “Social” category improved performance from 85.0% to 87.7%.

5.2 Processing the Complete Dataset

Based on the results outlined above, the full dataset of 60,647 entities was processed with the 7-prompt ensemble method, resulting in total of 424,529 LLM queries. The run was carried out on a cluster system using the vLLM library for efficient LLM generation. The system achieved 2219.38 tokens/s input and 54.05 tokens/s output. Despite identical formatting instructions, the prompts varied in reliability: the best required only 1.1 attempts per entity on average, while the most problematic needed 1.9 attempts (maximum 5 attempts per entity). Most failures stemmed from malformed JSON or empty responses; others involved incorrect category choices or hallucinated words. After majority voting across all 7 prompts, we obtained the final LLM annotation of all entities.

For Q1 (Categories), the most common labels were Administrative (19.5%, n=11,844), Professional/Work-related (17.9%, n=10,830), and General Social Group (12.8%, n=7,744), reflecting the formal organizational structure of post-war Finnish society. Multi-label classification proved essential for Q1: 29.2% of entities (17,717) received multiple category labels, with all 18 categories appearing in multi-label combinations (Table 5). Administrative entities most frequently combined with Professional/Work-related (1,654 cases), Educational/Academic (1,583), and Social Welfare (1,072). In contrast, Q2–Q4 showed minimal multi-label usage (1.8–6.6%), validating our design where Q1 captures overlapping domains while Q2–Q4 capture more singular dimensions. The full distribution across all labels is listed in Appendix A.4.

5.2.1 Mention-Weighted Analysis

All results so far were reported for unique entities. Weighting by mention frequency of each unique entity reveals distinct patterns between hobbies and organizations (full details are in Appendix A.5).

Table 5: Multi-label entity distribution across questions.

Question	Multi-label %	Total	Labels
Q1 – Category	29.2%	17,717	18
Q2 – Group size	6.6%	4,014	3
Q3 – Frequency	1.8%	1,076	4
Q4 – Physical act.	2.5%	1,511	4

Hobbies concentrate heavily in Creative/Artistic (36.3%) and Nature-related (32.9%) categories, while organizations distribute across General Social Group (19.0%), Administrative (12.6%), and Professional/Work-related (12.3%). These choices represent the final dataset and label distribution across the data. Group size (Q2) shows an interesting contrast whereby 96.1% of organizational mentions involve large groups, while hobbies split between solitary (43.4%) and small group (40.7%) activities. For frequency (Q3), 79.5% of organizational mentions are occasional, reflecting periodic meetings, whereas 66.0% of hobby mentions are continuous activities. Physical activity (Q4) distributions also diverge: 68.2% of organizations are stationary versus only 22.4% of hobbies, with most hobbies involving light (55.5%) or intense (10.0%) physical activity. These patterns validate our multi-dimensional approach and provide quantitative characterization of social participation patterns in the Karelian refugee population suitable for downstream health outcomes research.

5.3 Dimension Reliability by Category (Q1)

Any downstream analysis will need to take into account the reliability of both human and LLM predictions. For instance, Q4 (Physical Activity) performed clearly worst on the test set (Table 4). To investigate whether the prediction reliability depends on the entity’s category type, we analyzed Q2–Q4 reliability conditional on the answer to Q1. This addresses a question for downstream research: given that an entity is classified as a certain category (Q1), how reliable are answers to the other dimensions Q2–Q4? The statistics are shown in Table 6.

For instance: Sports/Physical activity entities (n=24), the model correctly identified the Q1 category 82% of the time. Among these sports-related entities, Q2 (Group Size) achieved 81% accuracy, Q3 (Frequency) 72%, and Q4 (Physical Activity) 76%. In contrast Creative/Artistic entities (n=33) achieved 93% Q1 accuracy but only 51% on Q4, reflecting difficulty distinguishing stationary activities from movement-based ones.

Q1 Category	n	Q1	Q2	Q3	Q4
Creative/Artistic	33	93/95	85/91	90/94	51/87
Professional	27	77/88	91/96	70/81	63/77
Sports/Physical	24	82/93	81/89	72/85	76/90
Administrative	21	85/88	95/96	76/79	84/83
Nature-related	19	79/92	75/86	60/71	58/72
General social	12	75/88	92/97	47/82	0/65
Educational	10	67/86	90/96	73/82	90/95
Property/Own.	10	67/69	73/82	63/85	80/72
Animal-related	7	83/94	79/86	81/82	71/73
Special hobbies	7	50/84	90/91	90/87	67/76
Military-related	6	91/91	83/92	44/78	0/79
Religious/Spirit.	6	80/92	86/99	67/87	67/78
Social welfare	5	59/67	100/100	40/76	20/78

Table 6: Q1–Q4 reliability conditioned by entity category in Q1, showing both model and human agreement rates. Format: Model% / Human%. Colors: $\geq 75\%$, 50–74%, $< 50\%$ (based on lower score).

Perfect agreement (0%)	19%
Low disag. (<25%)	44%
Moderate disag. (25–50%)	30%
High disag. (>50%)	7%

Table 7: Degrees of model–human disagreement.

6 Error analysis

As shown in Table 7, 63% of entities show perfect agreement or low disagreement between the model and human annotators. The full listing of entities and their human and model agreement levels is available in the Appendix A.6. We carry out a manual inspection of several main classes of entities by agreement:

Perfect agreement: Organizations with clear categories, group sizes, and physical activity levels are straightforward cases for the model. Examples: Choir (*Laulukuoro*), Bandy (*jääpallo*), Basketball and Volleyball (*kori- ja lentopalloilu*), Biathlon (*ampumahiihto*).

Model blind spots: Entities showing a high model disagreement despite high degree of human agreement. Typical cases are abbreviations without explicit context, and culturally-specific items such as rare Finnish folk dances. These are often known to humans familiar with Finnish culture and history, but unknown to general open-weight LLMs. Another typical case is the model’s overconfidence, making a guess, where humans resort to the “Cannot be determined” label. Typical examples are e.g. Seniors’ circle (*vanhusten piiri*), and Lotta Svärd

activities (*lottatoiminta*).

Genuine Ambiguity: Entities with high both model and human disagreement include e.g. Motor sports (*moottoriurheilu*), where the human annotators did not agree on frequency, group size, nor degree of physical activity. Another typical class are abbreviations, such as the military regiment abbreviation (*2.KKK/KKR*) and entities with OCR errors, where some annotators attempted disambiguation, while others marked a “Data error”.

7 Conclusions

We demonstrate that LLMs can implement social organization and leisure activity classification according to a custom taxonomy developed by domain experts for a specific analytic task. The models achieve 76.8% F-score, which is 93.9% of human baseline performance on coarse labels, and 93.8% on full granularity. This includes the added complexity of a relatively small language (Finnish) and the need of multi-label classification, where we see the LLMs are capable of assigning multiple labels, allowing for a structured multi-question schema, where some questions (e.g. Q1) are inherently multi-label.

The performance of the LLM at nearly 94% of human annotation provides a worthwhile alternative to manual classification of tens of thousands of entities, especially if any downstream analysis can be designed to take into account the levels of uncertainty necessarily present in such a categorization task.

To this end, we find that analyzing and quantifying LLM classification errors is crucial, since these errors are rarely random. They follow consistent patterns: some reflect genuine failures of model understanding, while others represent interpretive differences where the model’s reasoning diverges from human consensus rather than being outright wrong. Characterizing these error types, identifying which options are systematically misclassified, and providing reliability measures for specific categories is essential for downstream use of this data.

All data will be made available upon publication under an open license.

8 Acknowledgments

This work was carried out in the Human Diversity University profilation programme (PROFI-7) of the Research Council of Finland (grant 352727), and in part supported also through the KinSocieties,

ERC-2022-ADG grant number 101098266. Computational resources were provided by CSC — the Finnish IT Center for Science.

9 Limitations

LLMs show consistent failure patterns, including uncertainty avoidance and cultural knowledge gaps. Rather than selecting “Cannot be determined,” the models tend to guess concrete answers even when annotators acknowledge ambiguity. Similarly, Finnish-specific terminology and regional abbreviations lead to systematic misclassification.

Additionally developing a custom taxonomy requires substantial annotation effort. Our annotation framework involved five rounds of refinement with four annotators before achieving stable guidelines. Comparing this to established schemes like NER, where pretrained models and labeled datasets exist, this represents a significant overhead. Furthermore, if reliability analysis reveals poor alignment for specific categories, additional work is required, either more annotations to improve agreement, or refining the categories themselves. However while the framework is designed to address the specific research questions related to our data and may not be directly transferable to others, we believe that many of its categories are of universal interest (e.g. in health research), and that the framework can serve as a reference to be adapted for other works describing social participation.

References

- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. [Don’t count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 238–247, Baltimore, Maryland. Association for Computational Linguistics.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, pages 4171–4186.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The Llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. [Deepseek-r1 incentivizes reasoning in llms through reinforcement learning](#). *Nature*, 645(8081):633–638.
- Helga Hansdóttir, María K. Jonsdóttir, Diana E. Fisher, Gudny Eiríksdóttir, Pálmi V. Jonsson, and Vilundur Gudnason. 2022. [Creativity, leisure activities, social engagement and cognitive impairment: the AGES-Reykjavík study](#). *Aging Clinical and Experimental Research*, 34:1027–1035.
- Julianne Holt-Lunstad, Timothy B. Smith, and J. Bradley Layton. 2010. [Social relationships and mortality risk: A meta-analytic review](#). *PLoS Medicine*, 7(7):e1000316.
- Chen Huang and Guoxiu He. 2025. [Text clustering as classification with LLMs](#). *Preprint*, arXiv:2410.00927.
- Qianyu Huang and Tongfang Zhao. 2024. [Leveraging large language models for entity matching](#). *Preprint*, arXiv:2405.20624.
- Yoshi Iwasaki, Jiri Zuzanek, and Roger C. Mannell. 2001. [The effects of physically active leisure on stress-health relationships](#). *Canadian Journal of Public Health*, 92(3):214–218.
- Joonatan Laato, Jenna Kanerva, John Loehr, Virpi Lummaa, and Filip Ginter. 2024. [Extracting social connections from Finnish Karelian refugee interviews using LLMs](#). In *Proceedings of Fifth Conference on Computational Humanities Research (CHR’2024)*.
- Mengfei Lan, Lecheng Zheng, Shufan Ming, and Halil Kilicoglu. 2024. [Multi-label sequential sentence classification via large language model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16086–16104, Miami, Florida, USA. Association for Computational Linguistics.
- John Loehr, Robert Lynch, Johanna Mappes, Tuomas Salmi, Jenni Pettay, and Virpi Lummaa. 2017. [Newly digitized database reveals the lives and families of forced migrants from Finnish Karelia](#). *Nordic Yearbook of Population Research*, 52:59–70.
- Marcus Ma, Georgios Chochlakakis, Niyantha Maruthu Pandiyan, Jesse Thomason, and Shrikanth Narayanan. 2025. [Large language models do multi-label classification differently](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2472–2495, Suzhou, China. Association for Computational Linguistics.

- Laura McQuade and Roger O’Sullivan. 2024. [Examining arts and creativity in later life and its impact on older people’s health and wellbeing: a systematic review of the evidence](#). *Perspectives in Public Health*, 144(6):344–353.
- Meta-AI. 2025. Llama 4 Scout 17B–16E Model Card. <https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E>.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proceedings of ICLR’13*.
- Mistral-AI. 2024. Mistral-Large-Instruct-2407 Model Card. <https://huggingface.co/mistralai/Mistral-Large-Instruct-2407>.
- Ryan E. Rhodes, Ian Janssen, Shannon S. D. Bredin, Darren E. R. Warburton, and Adrian Bauman. 2017. [Physical activity: Health impact, prevalence, correlates and interventions](#). *Psychology & Health*, 32(8):942–975.
- Victor S. Sheng, Foster Provost, and Panagiotis G. Ipeirotis. 2008. [Get another label? improving data quality and data mining using multiple, noisy labelers](#). In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD ’08, page 614–622, New York, NY, USA. Association for Computing Machinery.
- Rion Snow, Brendan O’Connor, Daniel Jurafsky, and Andrew Ng. 2008. [Cheap and fast – but is it good? Evaluating non-expert annotations for natural language tasks](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 254–263, Honolulu, Hawaii. Association for Computational Linguistics.
- Seyed Amin Tabatabaei, Sarah Fancher, Michael Parsons, and Arian Askari. 2025. [Can large language models serve as effective classifiers for hierarchical multi-label classification of scientific documents at industrial scale?](#) In *Proceedings of the 31st International Conference on Computational Linguistics: Industry Track*, pages 163–174, Abu Dhabi, UAE. Association for Computational Linguistics.
- Vijay Viswanathan, Kiril Gashteovski, Carolin Lawrence, Tongshuang Wu, and Graham Neubig. 2024. [Large language models enable few-shot clustering](#). *Transactions of the Association for Computational Linguistics*, 12:321–333.
- Dorota Węziak-Białowolska, Piotr Białowolski, and Pier Luigi Sacco. 2023. [Mind-stimulating leisure activities: Prospective associations with health, wellbeing, and longevity](#). *Frontiers in Public Health*, 11:1117822.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 23 others. 2025b. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. 2024. [Large language models as optimizers](#). In *Proceedings of The Twelfth International Conference on Learning Representations*.
- Lauren A. Yates, Shier Ziser, Aimee Spector, and Martin Orrell. 2016. [Cognitive leisure activities and future risk of cognitive impairment and dementia: systematic review and meta-analysis](#). *International Psychogeriatrics*, 28(11):1791–1806.
- Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2023. [Large language models are human-level prompt engineers](#). In *Proceedings of The Eleventh International Conference on Learning Representations*.

A Appendix

A.1 Task prompt

TASK

Classify the entity given below by answering four questions (q1-q4).
For each question, choose the option(s) that **most accurately** match the entity, following the
↪ instructions provided.
Write your answers exactly according to the **CORRECT ANSWER FORMAT** structure given below.

QUESTIONS AND OPTIONS

Always clearly separate the option and its description, e.g.
`Professional/Work-related - Trade unions, professional associations, workplace-related groups. . .`

Question 1 (q1) - Categories (select **all** that apply)

Option	Description
Professional/Work-related	Trade unions, professional associations, workplace-related groups;
Cultural/Traditional	Cultural heritage, local culture, traditions (e.g. Karelian Society)
Religious/Spiritual	Church activities, religious education, spiritual associations
Sports/Physical activity	All kinds of sports, exercise, team sports and individual sports
Creative/Artistic	Music, visual arts, handicrafts, theatre, dance, etc.
Educational/Academic	Studying, education, scholarly associations, study circles
Social welfare	Charity work, volunteering, community service
Administrative	Public administration roles and bodies (municipality, church, state), ↪ non-political
Political	Parties, party organizations, political advocacy
General social group	Martta associations, youth clubs, general-purpose community groups
Nature-related	Hunting, fishing, hiking, nature conservation
Health-related	Peer-support groups, support for people with illnesses (e.g. Rheumatism ↪ Association)
Property/Ownership	Road maintenance associations, housing companies, detached-house maintenance ↪ associations
Non-physical games	Chess, board games, role-playing and video games
Animal care/Hobby	Beekeeping, horse riding, dog shows
Special hobbies	Stamp collecting, radio technology, etc.
Cooking	Food preparation, baking
Military-related	National defence, reservist activities
Not definable	Impossible to determine from name/context
Data error	Entity is not an organization/hobby OR is a data error

Question 2 (q2) - Group size (select the most typical one(s))

Option	Description
Alone	1 person; the activity can be done completely independently
Small group	2-5 people; close cooperation between participants
Large group	> 6 people; requires broader organization and structures
Not definable	Group size varies / cannot be inferred
Data error	Data issue prevents assessment

Question 3 (q3) - Frequency of activity (select the most likely one(s))

Option	Description
Regular	Fixed schedule, e.g. weekly
Occasional	About once a month, not fully fixed
Event-based	Once a year or less often, in connection with an event

```
| Continuous | Ongoing activity without fixed schedule; can be done any time |
| Not definable | Frequency is unclear |
| Data error | Data issue prevents assessment |
```

Question 4 (q4) – Level of movement / Physical activity (select the most likely one(s))

```
| Option | Description |
|-----|-----|
| Intense | Heavy sports, intensive muscular work |
| Continuous | Long-lasting steady movement (e.g. hiking) |
| Light | Occasional light movement, light activity/puttering |
| Stationary | Minimal movement, sitting/standing (meetings, administration) |
| Not definable | Physical activity level cannot be determined |
| Data error | Data issue prevents assessment |
```

ANNOTATION GUIDELINES

General principles

- Accuracy > coverage → if uncertain, choose **“Not definable”**.
- If the entity is **not** an organization/hobby, use **“Data error”** for all questions.
- Base decisions on the available information; avoid guessing.
- If provided, take into account the hierarchy in which the organization is mentioned. If it is empty → or UNK, you can ignore it.

Question-specific notes

q1: Select at least one category; some entities belong to multiple categories.
* Administrative is not Political: administrative = official governing bodies, political = parties.

q2: Choose the most typical group size; meetings/organizations → often **Large group**.
* If the estimate is difficult → **Not definable**.

q3:
* **Continuous** = independent, no group required.
* **Regular** = active weekly/monthly meetings.
* **Occasional** = committees & boards (~ once/month).
* **Event-based** = rarer, 1-2 times/year.
* Unclear → **Not definable**.

q4:
* Administrative body → **Stationary**; sports club → **Intense**.
* Board of a sports club → **Stationary**.
* Unclear → **Not definable**.

EXAMPLES

```
|-----|----|----|----|----|
| football | Sports/Physical activity | Large group | Regular | Intense |
| rheumatism association | Health-related | Large group | Regular | Stationary |
| mushroom picking | Nature-related | Alone, Small group | Continuous | Light |
| social welfare committee | Administrative, Social welfare | Large group | Occasional | Stationary |
| Guards band | Military-related, Creative/Artistic | Large group | Regular | Stationary |
| Järvenpää youth board | Administrative | Large group | Occasional | Stationary |
| housekeeping | Property/Ownership | Alone, Small group | Continuous | Light |
| local party branch | Political | Large group | Occasional | Stationary |
```

CORRECT ANSWER FORMAT

[Answer begin]

```
{{
  "{{entity_name}}": {{
    "q1": [...],
    "q2": [...],
    "q3": [...],
    "q4": [...]
  }}
}}
```

```

}}
[Answer end]

#### Entity to be annotated

"{entity_name}"

Hierarchy: "{hierarchies}"

"{past_mistakes}"

```

A.2 Inference Hyperparameters

Inference was performed on the LUMI supercomputer using vLLM as the inference engine with the following sampling parameters: temperature=0.3, top_p=1.0, top_k=40, max_tokens=300. Processing used batch size of 150.

A.3 Inter-annotator agreement details

Annotator Pair	Q1	Q2	Q3	Q4	Average
Pairwise Cohen's Kappa – Evaluation Set (n=50)					
Annotator 1 vs 2	0.800	0.806	0.667	0.721	0.749
Annotator 1 vs 3	0.753	0.743	0.725	0.711	0.733
Annotator 1 vs 4	0.753	0.767	0.524	0.660	0.676
Annotator 2 vs 3	0.684	0.753	0.677	0.688	0.701
Annotator 2 vs 4	0.697	0.848	0.651	0.754	0.737
Annotator 3 vs 4	0.683	0.714	0.571	0.535	0.626
Mean	0.728	0.772	0.636	0.678	0.704
Pairwise Cohen's Kappa – Test Set (n=150)					
Annotator 1 vs 2	0.783	0.825	0.642	0.620	0.718
Annotator 1 vs 3	0.824	0.790	0.698	0.528	0.710
Annotator 1 vs 4	0.778	0.776	0.581	0.723	0.714
Annotator 2 vs 3	0.737	0.751	0.636	0.507	0.658
Annotator 2 vs 4	0.712	0.812	0.649	0.680	0.713
Annotator 3 vs 4	0.767	0.732	0.602	0.415	0.629
Mean	0.767	0.781	0.634	0.579	0.690
Individual F1 vs Leave-One-Out Majority – Evaluation Set (n=50)					
Annotator 1	0.867	0.873	0.813	0.807	0.840
Annotator 2	0.793	0.907	0.833	0.793	0.832
Annotator 3	0.756	0.813	0.833	0.733	0.784
Annotator 4	0.793	0.880	0.710	0.713	0.774
Mean	0.802	0.868	0.797	0.762	0.807
Individual F1 vs Leave-One-Out Majority – Test Set (n=150)					
Annotator 1	0.865	0.898	0.749	0.744	0.814
Annotator 2	0.800	0.902	0.756	0.722	0.795
Annotator 3	0.830	0.869	0.753	0.580	0.758
Annotator 4	0.805	0.876	0.696	0.724	0.775
Mean	0.825	0.886	0.738	0.693	0.786

Note: Krippendorff's Alpha values are nearly identical to Cohen's Kappa (within 0.001) and omitted for brevity.

Table 8: Inter-annotator agreement metrics across evaluation (n=50) and test (n=150) sets.

A.4 Full distribution of unique entities

Table 9: Distribution across 60,647 unique entities for all four questions.

Q1 – Category			Q2 – Group Size		
Label	n	%	Label	n	%
Administrative	11,844	19.5	Large group	53,574	88.3
Professional	10,830	17.9	Small group	4,889	8.1
General Social	7,744	12.8	Alone	4,840	8.0
Creative	7,396	12.2	Data error	853	1.4
Sports	7,328	12.1	Not definable	654	1.1
Social Welfare	5,760	9.5			
Educational	4,921	8.1			
Cultural	3,846	6.3			
Military	3,798	6.3			
Religious	3,166	5.2			
Property	2,696	4.4			
Nature	2,590	4.3			
Political	2,034	3.4			
Health	1,930	3.2			
Special hobbies	958	1.6			
Animal care	886	1.5			
Data error	853	1.4			
Not definable	283	0.5			
Cooking	176	0.3			
Non-phys. games	175	0.3			
Q3 – Frequency			Q4 – Physical		
Label	n	%	Label	n	%
Occasional	42,999	70.9	Stationary	41,103	67.8
Regular	11,800	19.5	Light	9,902	16.3
Continuous	5,092	8.4	Intense	6,291	10.4
Data error	853	1.4	Not definable	2,567	4.2
Event-based	555	0.9	Continuous	1,542	2.5
Not definable	490	0.8	Data error	853	1.4

A.5 Full distribution of entity occurrences

Label	Hobbies		Organizations		Total	
	Count	%	Count	%	Count	%
Q2 – Group Size						
Large group	55,252	12.9	198,318	96.1	253,570	39.9
Alone	186,679	43.4	512	0.2	187,191	29.4
Small group	174,877	40.7	4,104	2.0	178,981	28.1
Not definable	12,232	2.8	1,140	0.6	13,372	2.1
Data error	687	0.2	2,387	1.2	3,074	0.5
Q3 – Frequency						
Occasional	43,018	16.0	163,126	79.5	206,144	43.4
Continuous	177,956	66.0	1,804	0.9	179,760	37.8
Regular	41,123	15.2	36,439	17.8	77,562	16.3
Not definable	6,025	2.2	1,060	0.5	7,085	1.5
Data error	687	0.3	2,387	1.2	3,074	0.6
Event-based	865	0.3	448	0.2	1,313	0.3
Q4 – Physical Activity						
Stationary	59,508	22.4	145,285	68.2	204,793	42.7
Light	147,726	55.5	31,210	14.6	178,936	37.3
Intense	26,605	10.0	16,138	7.6	42,743	8.9
Not definable	11,513	4.3	14,497	6.8	26,010	5.4
Continuous	19,898	7.5	3,663	1.7	23,561	4.9
Data error	687	0.3	2,387	1.1	3,074	0.6

Table 10: Mention-weighted distribution (355,648 total mentions) for Q2–Q4 by entity type.

Category	Hobbies		Organizations		Total	
	Count	%	Count	%	Count	%
Creative/Artistic	97,180	36.3	16,639	6.3	113,819	21.4
Nature-related	88,174	32.9	3,880	1.5	92,054	17.3
Sports/Physical	46,269	17.3	15,390	5.8	61,659	11.6
General Social Group	1,621	0.6	50,089	19.0	51,710	9.7
Professional/Work	1,811	0.7	32,344	12.3	34,155	6.4
Administrative	235	0.1	33,301	12.6	33,536	6.3
Cultural/Traditional	565	0.2	31,685	12.0	32,250	6.1
Educational/Academic	14,372	5.4	12,112	4.6	26,484	5.0
Social Welfare	1,055	0.4	21,406	8.1	22,461	4.2
Religious/Spiritual	2,678	1.0	12,440	4.7	15,118	2.8
Military	433	0.2	13,858	5.3	14,291	2.7
Property/Ownership	3,931	1.5	5,617	2.1	9,548	1.8
Health-related	136	0.1	5,482	2.1	5,618	1.1
Animal care/hobby	3,689	1.4	967	0.4	4,656	0.9
Political	220	0.1	4,216	1.6	4,436	0.8
Special hobbies	3,239	1.2	799	0.3	4,038	0.8
Data error	687	0.3	2,387	0.9	3,074	0.6
Not definable	278	0.1	918	0.3	1,196	0.2
Non-physical games	880	0.3	211	0.1	1,091	0.2
Cooking	565	0.2	28	0.0	593	0.1

Table 11: Q1 category distribution weighted by mention frequency (355,648 total mentions).

A.6 Full test set categorization

Rank	Entity	Model Disagr	Human Disagr
1	Loimaa Co-op Bank (<i>Loimaan KOP</i>)	100.0%	66.7%
2	Folk dance (<i>tanhu</i>)	100.0%	50.7%
3	Poor relief fund (<i>Vaivaiskassa</i>)	100.0%	66.7%
4	Military regiment abbr. (2.KKK/KKR)	91.7%	75.0%
5	Association abbr. (<i>KTV</i>)	87.5%	25.0%
6	Summer cottage life (<i>kesämökkielämä</i>)	79.2%	56.2%
7	Sea excursions (<i>retkeily merellä</i>)	79.2%	54.2%
8	Boating sports (<i>venueurheilu</i>)	79.2%	65.3%
9	Nokia SOK evening club (<i>Nokian SOK:n ilkkakerho</i>)	75.0%	50.0%
10	Council (<i>valtuusto</i>)	75.0%	68.8%
11	Seniors' circle (<i>vanhusten piiri</i>)	75.0%	18.8%
12	Self-employed entrepreneurs (<i>Yksityisyrittäjät</i>)	75.0%	20.8%
13	Ylöjärvi TPSL branch (<i>Ylöjärven TPSL</i>)	75.0%	27.8%
14	Staying at summer cottage (<i>kesämökillä oleskelu</i>)	66.7%	52.1%
15	Summer cottage stays (<i>kesämökkeily</i>)	66.7%	52.1%
16	Motor sports (<i>moottoriurheilu</i>)	66.7%	75.0%
17	Camping/tents (<i>telttailu</i>)	66.7%	25.0%
18	Karkkula Farmers' Association (<i>Karkkulan Maamiesseura</i>)	62.5%	31.9%
19	Bird banding (<i>lintujen merkitsijänä toimiminen</i>)	62.5%	41.7%
20	Martha tasks (<i>martha-työt</i>)	62.5%	29.2%
21	Martha organization work (<i>marttatyö</i>)	62.5%	29.2%
22	Attending sewing circles (<i>ompeluseuroissa käyminen</i>)	54.2%	31.2%
23	War veterans (<i>Sotaveteraanit</i>)	54.2%	27.1%
24	Hauho Farmers' Association (<i>Hauhon Maamiesseura</i>)	50.0%	29.9%
25	Hämeenlinna Camera Club (<i>Hämeenlinnan Kameraseura</i>)	50.0%	29.2%
26	Dev. disabilities support (<i>Kehitysvammaisten tukiyhdistys</i>)	50.0%	29.9%
27	Lotta Svärd organization (<i>Lotta-Svärd-järjestö</i>)	50.0%	20.8%
28	Lotta Svärd activities (<i>lottatoiminta</i>)	50.0%	20.8%
29	Martha activities (<i>marttatoiminta</i>)	50.0%	25.0%
30	Travels (<i>matkusteleee</i>)	50.0%	66.7%
31	Opera (<i>ooppera</i>)	50.0%	56.2%
32	Local horse association (<i>paikallinen hevosyhdistys</i>)	50.0%	52.8%
33	Gliding (<i>purjelento</i>)	50.0%	64.6%
34	Chess (<i>shakki</i>)	50.0%	20.8%
35	Betting (<i>veikkaus</i>)	50.0%	58.3%
36	Diaconia Committee (<i>Diakoniatuomikunta</i>)	45.8%	34.7%
37	Workers' association (<i>työväenyhdistys</i>)	45.8%	41.7%
38	Church attendance (<i>kirkossakäyminen</i>)	41.7%	0.0%
39	East Karelia Folk College (<i>Itä-Karjalan kansanopisto</i>)	37.5%	33.3%
40	Making canal boards (<i>kanavataulujen teko</i>)	37.5%	0.0%
41	Middle-distance running (<i>keskipitkien matkojen juoksu</i>)	37.5%	0.0%
42	Fitness gymnastics (<i>kuntovoimistelu</i>)	37.5%	33.3%
43	Handicrafts (<i>käsityö</i>)	37.5%	6.2%
44	Agricultural Producers (<i>Maataloustuottajat</i>)	37.5%	27.1%
45	Martha Association (<i>Martta-yhdistys</i>)	37.5%	27.1%
46	Mosaic work (<i>mosaiikkityöt</i>)	37.5%	6.2%
47	Archaeological artifacts (<i>muinaistieteellisten esineiden kerääminen</i>)	37.5%	45.8%
48	Drawing (<i>piirustus</i>)	37.5%	0.0%
49	Ice fishing (<i>pilkkiönginta</i>)	37.5%	16.7%
50	Rotary Club activity (<i>rotarytoiminta</i>)	37.5%	35.4%
51	Selkämeri NCO Women (<i>Selkämeren Alipäällystönaiset ry</i>)	37.5%	25.0%
52	Playing chess (<i>shakinpeluu</i>)	37.5%	25.0%
53	Folk dances (<i>tanhut</i>)	37.5%	45.8%
54	Täkänä weaving (<i>täkänäin kudonta</i>)	37.5%	18.8%
55	Viipuri Guild (<i>Viipurin Kilta ry</i>)	37.5%	25.0%
56	Visuvesi Workers' Assoc. (<i>Visuvesen Työväenyhdistys</i>)	37.5%	50.7%
57	Horse breeding (<i>hevoskasvatus</i>)	29.2%	29.2%
58	Baltic Sea Summer (<i>Itämeren Kesä</i>)	29.2%	36.8%
59	Agrarian League (<i>Maalaisliitto</i>)	29.2%	36.8%
60	Professional Drivers' Union (<i>Ammattiautoilijaliitto</i>)	25.0%	27.1%
61	Balalaika playing (<i>balalaikan soittaminen</i>)	25.0%	0.0%
62	Wilderness Scouts (<i>Eränkävijät</i>)	25.0%	50.0%
63	Heinola Rural Road Board (<i>Heinolan mlk:n tienhoitokunta</i>)	25.0%	33.3%
64	Skiing (<i>Hiihto</i>)	25.0%	12.5%
65	Janakkala Society (<i>Janakkala-seura</i>)	25.0%	47.9%
66	Fabric weaving (<i>Jcankaankutominen</i>)	25.0%	75.0%
67	Cabbage cultivation (<i>kaalinviljely</i>)	25.0%	27.1%
68	Property Inspection Board (<i>Kiinteistö tarkastuslautakunta</i>)	25.0%	41.7%
69	Domestic chores (<i>kodintyöt</i>)	25.0%	41.7%
70	Making 3D pictures (<i>kolmiulotteisten taulujen tekeminen</i>)	25.0%	12.5%
71	Typing (<i>konekirjoitus</i>)	25.0%	12.5%
72	Household chores (<i>kotitaloustyöt</i>)	25.0%	20.8%
73	Jogging (<i>lenkkeily</i>)	25.0%	6.2%
74	Salmon farming (<i>lohen kasvatus</i>)	25.0%	62.5%
75	Casual reading (<i>lueskeleminen</i>)	25.0%	0.0%
76	Farmers' Assoc. sewing (<i>Maamiesseuran ompeluseurassa käynti</i>)	25.0%	41.0%
77	Rag rug weaving (<i>matonkutominen</i>)	25.0%	18.8%
78	Beekeeping (<i>mehiläisten hoito</i>)	25.0%	8.3%
79	Study circle (<i>Opintokerho</i>)	25.0%	27.1%
80	Lace crocheting (<i>pitsinvirkkaus</i>)	25.0%	12.5%
81	Norden Association (<i>Pohjola-Norden</i>)	25.0%	16.7%
82	Polio Disabled Association (<i>Polioinvalidit ry</i>)	25.0%	12.5%
83	Women of Foremen's Assoc. (<i>Rakennusmestariyhdistyksen Naisjaosto</i>)	25.0%	22.9%

Rank	Entity	Model Disagr	Human Disagr
84	Race walking (<i>ratakävely</i>)	25.0%	45.8%
85	Game management (<i>riistanhoito</i>)	25.0%	47.9%
86	Making ryijy rugs (<i>ryijyjen tekeminen</i>)	25.0%	22.9%
87	Rambling/hiking (<i>samoilu</i>)	25.0%	31.2%
88	Knitting socks (<i>sukkien neulominen</i>)	25.0%	0.0%
89	Finnish Real Estate Federation (<i>Suomen Kiinteistöliitto</i>)	25.0%	41.7%
90	Finland–Soviet Union Society (<i>Suomi-Neuvostoliitto-Seura</i>)	25.0%	54.2%
91	Studying stars/astromony (<i>tutkii tähtiä</i>)	25.0%	18.8%
92	Workers' drama club (<i>Työväenyhdistyksen näytelmäkerho</i>)	25.0%	18.8%
93	Vaasa Reserve Officers' Club (<i>Vaasan Reserviupseerikerho</i>)	25.0%	27.1%
94	Sewing clothes (<i>vaatteiden ompelu</i>)	25.0%	6.2%
95	Carving course (<i>veistokurssi</i>)	25.0%	43.8%
96	Viipuri Workers' Institute (<i>Viipurin Työväenopisto</i>)	25.0%	37.5%
97	Violin playing (<i>viulun soitto</i>)	25.0%	12.5%
98	Imatra Region Co-op Bank (<i>Imatran seudun Osuuspankki</i>)	16.7%	34.0%
99	Agricultural Club Association (<i>Maatalouskerhoyhdistys</i>)	16.7%	38.2%
100	Furniture making (<i>huonekalujen valmistus</i>)	12.5%	12.5%
101	Poultry keeping (<i>kananhoito</i>)	12.5%	16.7%
102	Karelia Athletes (<i>Karjalan Urheilijat</i>)	12.5%	8.3%
103	Javelin throw (<i>keihäänheitto</i>)	12.5%	0.0%
104	Home gardening (<i>kotipuutarhanhoito</i>)	12.5%	27.1%
105	Kuusaa Horticultural Assoc. (<i>Kuusaa Puutarhayhdistys</i>)	12.5%	40.3%
106	Ship Officers' Association (<i>Laivanpäällystöliitto</i>)	12.5%	33.3%
107	Lappeenrannan evankelisluterilaisen seurakunta	12.5%	33.3%
108	Land Acquisition Board (<i>maanlunastuslautakunta</i>)	12.5%	39.6%
109	Seamen's Mission (<i>merimieslähetys</i>)	12.5%	49.7%
110	Metalworkers' Union (<i>Metallityöväen Liitto</i>)	12.5%	20.8%
111	Bankers' Union (<i>Pankkimiesliitto ry</i>)	12.5%	33.3%
112	Paper Workers' Guild (<i>Paperiammattikunta</i>)	12.5%	25.0%
113	Philatelic club (<i>Postimerkkikerho</i>)	12.5%	20.8%
114	Carpentry workshop (<i>puusepänerastas</i>)	12.5%	37.5%
115	Construction Workers' Union (<i>Rakennustyöväen Liitto</i>)	12.5%	27.1%
116	Railway Workers' Choir (<i>rautatieläisten kuoro</i>)	12.5%	0.0%
117	Railway Workers' Union (<i>Rautatieläisten liitto</i>)	12.5%	27.1%
118	Restaurant Staff Union (<i>Ravintolahenkilökunnan Liitto</i>)	12.5%	27.1%
119	Taxation Board (<i>taksoituslautakunta</i>)	12.5%	22.9%
120	Playing tennis (<i>tenniksen pelaaminen</i>)	12.5%	14.6%
121	VR Hyvinkää Men's Choir (<i>VR:n Hyvinkään mieskuoro</i>)	12.5%	13.9%
122	Biathlon (<i>ampumahiihto</i>)	0.0%	14.6%
123	Church Council of Haapavesi (<i>Haapaveden kirkkoneuvosto</i>)	0.0%	0.0%
124	Hämeen Tarmo Women Gymnasts (<i>Hämeen Tarmon Naisvoimistelijat</i>)	0.0%	6.2%
125	Imatra Orienteers (<i>Imatran Suunnistajat</i>)	0.0%	12.5%
126	Bandy (<i>jääpallo</i>)	0.0%	0.0%
127	Primary School Board (<i>Kansakoululautakunta</i>)	0.0%	20.1%
128	Primary school board (<i>kansakoululautakunta</i>)	0.0%	18.8%
129	Kemi Rural Labour Committee (<i>Kemin mlk:n työvoimatoimikunta</i>)	0.0%	25.0%
130	Basketball and Volleyball (<i>kori- ja lentopalloilu</i>)	0.0%	0.0%
131	Municipal Council (<i>Kunnanvaltuusto</i>)	0.0%	18.8%
132	Assoc. Municipal Officials (<i>Kuntayhtymän virkamiesyhdistys ry</i>)	0.0%	31.2%
133	Singing Masters (<i>Laulavat mestarit</i>)	0.0%	0.0%
134	Choir (<i>Laulukuoro</i>)	0.0%	0.0%
135	Lokalaiti Men's Choir (<i>Lokalahden Laulumiehet</i>)	0.0%	0.0%
136	Forestry (<i>metsänhoito</i>)	0.0%	31.2%
137	Men's choir (<i>mieslaulajiin</i>)	0.0%	0.0%
138	Women Gymnasts (<i>Naisvoimistelijat</i>)	0.0%	6.2%
139	Drama club activity (<i>näytelmäkerhotoiminta</i>)	0.0%	12.5%
140	Stage activities/theatre (<i>näyttämötoiminta</i>)	0.0%	12.5%
141	Ball sports (<i>palloilu</i>)	0.0%	12.5%
142	Pukkila Municipal Council (<i>Pukkilan kunnanvaltuusto</i>)	0.0%	12.5%
143	Radio listening (<i>radion kuuntelu</i>)	0.0%	0.0%
144	Mixed choir (<i>sekakuoro</i>)	0.0%	0.0%
145	Non-fiction literature (<i>tietopuolinen kirjallisuus</i>)	0.0%	6.2%
146	Workers' Athletes (<i>Työväen Urheilijat</i>)	0.0%	0.0%
147	Vakka Transport (<i>Vakka-Kuljetus</i>)	0.0%	54.2%
148	Boat building (<i>veneiden teko</i>)	0.0%	18.8%
149	Studying Russian (<i>venäjänkielen opiskelu</i>)	0.0%	22.9%
150	Viljakkala Sports Club (<i>Viljakkalan urheiluseura</i>)	0.0%	0.0%

Table 12: Model and human disagreement over the 150 test set entities.