



HAL
open science

Collaborative systems for large-scale online citizen participation

William Aboucaya

► **To cite this version:**

William Aboucaya. Collaborative systems for large-scale online citizen participation. Computers and Society [cs.CY]. Sorbonne Université, 2023. English. NNT : 2023SORUS461 . tel-04401421

HAL Id: tel-04401421

<https://inria.hal.science/tel-04401421>

Submitted on 17 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

SORBONNE UNIVERSITÉ

ÉCOLE DOCTORALE 130

INRIA Paris – Équipe projet MiMove

THÈSE

pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ SORBONNE UNIVERSITÉ

Discipline : Informatique

Présentée et soutenue par :

William ABOUCAYA

le : 30 novembre 2023

**Collaborative systems for large-scale online
citizen participation**

Sous la direction de :

M. Nikolaos GEORGANTAS – Directeur de thèse, INRIA Paris

M. Rafael ANGARITA – Co-encadrant, Université Paris Nanterre | LIP6 Sorbonne Université

Mme Valérie ISSARNY – Ancienne directrice de thèse, INRIA Paris

Rapporteurs :

Mme Claudia IGNAT – Chargée de Recherche HDR, INRIA Nancy-Grand Est

M. Christophe PRIEUR - Professeur, LISIS Université Gustave Eiffel

Présidente du jury :

Mme Vanda LUENGO – Professeure, LIP6 Sorbonne Université

Membres du jury :

Mme Sihem AMER-YAHIA - Professeure, CNRS | Université Grenoble Alpes

M. Iván CANTADOR – Associate Professor, Universidad Autónoma de Madrid

M. Nicolas MAUDET – Professeur, LIP6 Sorbonne Université

Abstract

Online participatory platforms have become a common means to involve citizens in public decision-making, allowing for participation at a larger scale than their offline counterparts, both in the number of participants and in the geographical distribution. However, the term “participatory platform” covers a wide range of extremely different systems, implying differences in the problems encountered by platforms administrators and contributors. More precisely, such platforms face specific issues when they aim at allowing citizens to collaborate to produce common contributions or when the number of contributors involved becomes particularly high. This Ph.D. research aims at identifying issues in contemporary online citizen participation platforms and proposing technical means to create participatory platforms more collaborative and suitable for large-scale online participation. My thesis is mainly based on previous works produced in the Computer-Supported Collaborative Work (CSCW) and Natural Language Processing (NLP) fields of computer science research.

The contributions of this thesis are: the identification of flaws in a specific citizen participation platform and the recommendation of platform design-oriented alternatives to solve them; the representation of a participatory platform as a knowledge graph and its enrichment using a preexisting external knowledge base; the identification of the different objectives motivating the creation of participatory platforms and of the different types of features for interaction implemented based on a series of interviews; the conception and implementation of a Natural Language Inference-based method to reduce issues faced by online citizen participation when the number of contributors becomes particularly high.

Résumé

Les plates-formes participatives en ligne sont devenues un moyen courant d'impliquer les citoyens dans la prise de décision publique, permettant une participation à plus grande échelle que leurs homologues hors ligne, à la fois en termes de nombre de participants et de répartition géographique. Cependant, le terme "plateforme participative" recouvre un large éventail de systèmes extrêmement différents, ce qui implique des différences dans les problèmes rencontrés par les administrateurs et les contributeurs des plateformes. Plus précisément, ces plateformes font face à des problèmes spécifiques lorsqu'elles visent à permettre aux citoyens de collaborer pour produire des contributions communes ou lorsque le nombre de contributeurs impliqués devient particulièrement élevé. Cette recherche doctorale vise à identifier les problèmes des plateformes de participation citoyenne contemporaines et à proposer des moyens techniques pour créer des plateformes participatives plus collaboratives et adaptées à une participation à grande échelle. Ma thèse s'appuie principalement sur des travaux antérieurs réalisés dans les champs de la recherche en informatique que sont le travail collaboratif assisté par ordinateur (CSCW) et le traitement du langage naturel (NLP).

Les contributions de cette thèse sont : l'identification des biais d'une plateforme participative spécifique et la recommandation d'alternatives centrées sur la conception de la plateforme pour les résoudre ; la représentation d'une plateforme participative sous la forme d'un graphe de connaissances (*knowledge graph*) et son enrichissement à partir d'une base de connaissances externe préexistantes ; l'identification des différents objectifs motivant la création de plateformes participatives et des différents types de fonctionnalités d'interaction mises en œuvre à partir d'une série d'entretiens ; la conception et la mise en œuvre d'une méthode basée sur l'inférence en langage naturel pour réduire les problèmes rencontrés par la participation citoyenne en ligne lorsque le nombre de contributeurs devient particulièrement élevé.

Acknowledgments

The last years of my Ph.D. thesis have come to an end. They have been an occasion for me to focus on a fascinating topic and I feel deeply grateful for everything I learned both academically and non academically. These years have also been an occasion for me to meet wonderful people without whom this work wouldn't be the same. First, I want to thank my Ph.D. supervisor Rafael Angarita for his support and guidance during these years. Thank you for giving me the opportunity to work with you and for every time our meetings have put me back on the right track.

Then, I want to thank my Ph.D. advisor Nikolaos Georgantas who accepted to supervise the last year of my doctorate. Thank you for welcoming me to the MiMove team and for giving me helpful advice during the last four years. I also want to thank the researchers Claudia Ignat and Christophe Prieur for accepting to review my thesis, and Sihem Amer-Yahia, Iván Cantador, Vanda Luengo and Nicolas Maudet for their participation in the jury for my Ph.D. defense. This acknowledgment extends to Julia Lawall, who accepted to be part of my annual evaluation committee.

Many thanks to my former and current colleagues in the MiMove team, Shahin Abdoul Soukour, Maroua Bahri, Zakaria Benomar, Lior Diler, Yifan Du, Nathalie Gaudechoux, Meriem Guemair, Diana Marino Duarte, Patient Ntumba, Françoise Sailhan and Haidong Zhao. Thank you for your support and for bringing joy onto our floor. I also want to thank Oana Balalau and Sonia Guehis with whom I had the chance to work on recent research works. Thank you for your great advice, your patience and for the energy you brought to our weekly meetings.

This work wouldn't have been possible without the support and love of my wonderful family, my mother Cécile, my sister Ilana and my late father Georges. Thank you for giving me the chance of being where I stand today and for filling me with joy and hope. I also want to thank my fantastic friends without whom my life during these four years and before wouldn't have been the same. I would like to express my heartfelt gratitude to Stanley Duplan and Prune Krotoff, whose presence when my doctoral project was accepted filled me with happiness, Anne Labatut for her continuous support during these years, and Morgane Leclerc, Camille Parrau and Hugo Plassais for giving me the confidence to apply for a Ph.D. I also thank my friends Fanny Eberlé, Enzo Gaccio and Elsa Goudenege for their support and encouragements during these four years and before.

Finally, I want to express my sincere gratitude to my first Ph.D. advisor Valérie Issary who supervised the first three years of my doctorate. Thank you for everything you taught me, for your guidance and for your support during these three years. Most importantly, thank you for your energy, for your humor and for your determination. It has been a true honor and pleasure to work with you.

Résumé étendu

1 Introduction

Ces dernières années, la participation citoyenne en ligne, ainsi que d'autres sujets inclus dans le champ des technologies civiques (*civic tech*), est devenue de plus en plus fréquemment utilisée avec pour but d'améliorer l'intégration des citoyens dans la prise de décisions publiques. La participation citoyenne en ligne peut prendre des formes diverses, telles que le budget participatif [71, 128], la collecte et la publication de données ouvertes citoyennes [28, 94] ou la consultation citoyenne [83, 148]. De plus, des plateformes participatives ont déjà été utilisées pour de nombreux sujets, par exemple la planification du développement de l'espace public [176], la transformation numérique d'un pays [91] ou la lutte et l'adaptation au changement climatique [63, 133]. Cependant, nous considérons que deux constantes peuvent être identifiées : une plateforme en ligne sur laquelle les citoyens sont amenés à produire et publier des contributions, et une entité organisatrice (par exemple, une mairie, une université) que nous appelons ici "institution" et qui initie le processus participatif. Dans cette thèse, ces éléments sont considérés comme constitutifs de la participation citoyenne en ligne.

L'objectif de cette thèse est de résoudre des problèmes auxquels sont confrontés les citoyens et les administrateurs de plateformes dans le contexte de la participation citoyenne en ligne. Certains de ces problèmes ont été identifiés précédemment par d'autres chercheurs [38, 75, 132] et d'autres sont identifiés dans les Chapitres 3 et 4 de cette thèse. Plus spécifiquement, nous cherchons à réduire la surcharge d'information à laquelle sont confrontés à la fois les citoyens et les administrateurs de plateformes, ainsi qu'à faciliter la collaboration entre les citoyens dans un contexte de participation en ligne à grande échelle. Nous nous concentrons sur la participation citoyenne à grande échelle parce que nous considérons que la scalabilité est l'un des principaux avantages de la participation en ligne par rapport à la participation en personne. Nous pouvons citer l'exemple que donne Robert Dahl en 1998, relevant le fait que donner à chaque citoyen d'une ville de 10.000 citoyens 10 minutes d'expression publique demanderait 200 jours à un rythme de 8 heures de discussion par jour [48]. Bien que ce problème soit particulièrement pertinent pour la participation en personne, il est largement réduit par la possibilité offerte à chacun d'exprimer ou de lire l'expression des autres à n'importe quel moment et à n'importe quel endroit. Par conséquent, les plateformes en ligne permettent à la participation citoyenne d'exister à de multiples échelles, depuis un conseil de quartier [146] à des initiatives impliquant plusieurs pays [20].

Pour répondre aux problèmes identifiés, nos contributions se basent sur des travaux pré-existants dans les champs de recherche en Interaction Humain-Machine (IHM), Travail Collaboratif Assisté par Ordinateur (TCAO) et Traitement Automa-

tique du Langage (TAL). Nous proposons de répondre aux deux questions suivantes :

1. Quels sont les problèmes rencontrés par les citoyens et les administrateurs de plateformes dans le contexte de la participation citoyenne en ligne ?¹
2. À partir de l'état actuel des connaissances en informatique, comment pouvons-nous réduire les problèmes identifiés afin de créer des plateformes participatives viables à grande échelle et permettant davantage de collaboration ?

1.1 Contributions

Les contributions de cette thèse sont concentrées autour de trois approches différentes pour améliorer la collaboration dans le cadre de la participation citoyenne à grande échelle. Nous dédions un chapitre de cette thèse à chacune de ces contributions. Leurs objectifs sont d'identifier ou de résoudre divers problèmes rencontrés par la participation citoyenne à grande échelle en s'appuyant sur différents champs de la recherche en informatique.

Étude d'une plateforme participative existante La première contribution de cette thèse est une étude d'une consultation citoyenne spécifique autour du projet de loi *République Numérique*. Nous étudions les interactions entre les contributions et les contributeurs de cette consultation afin d'identifier des failles dans le *design* de la plateforme. Nous proposons des recommandations de modifications de l'interface utilisateur et du *design* de la plateforme afin de réduire les biais engendrés par les failles identifiées.

De plus, nous proposons également une représentation des consultations citoyennes en ligne basée sur un graphe de connaissances. Cette représentation nous permet d'associer les propositions soumises dans une consultation à différents sujets et ressources issues de la base de connaissances DBpedia. Nous appliquons cette représentation à la consultation *République Numérique* afin d'améliorer les possibilités d'identification de contributions pertinentes pour un citoyen et d'aider les administrateurs de plateformes dans leur travail de synthèse des contributions.

La majorité de cette contribution a été publiée dans l'article suivant : William Aboucaya, Rafael Angarita, and Valérie Issarny. "Detecting Obstacles to Collaboration in an Online Participatory Democracy Platform: A Use-case Driven Analysis". In: 44th International Conference on Software Engineering (ICSE'22). Proceedings of the 2nd International Workshop on Equitable Data and Technology. FairWare '22. Pittsburgh, Pennsylvania: Association for Computing Machinery, May 2022, pp. 25–33 [1]. La partie de ce chapitre liée aux graphes de connaissances a été publiée séparément, dans l'article suivant : William Aboucaya, Sonia Guehis, and Rafael Angarita. "Building Online Public Consultation Knowledge Graphs". In: Text2KG 2023: International Workshop on Knowledge Graph Generation from Text, Collocated with the ESWC 2023. Ed. by CEUR-WS. Hersonissos (Crete), Greece, May 2023, p. 12. url: <https://inria.hal.science/hal-04076743>. [2].

¹Nous avons conscience que cette question a déjà été traitée par d'autres chercheurs avant nous, notre objectif est d'identifier de nouveaux problèmes.

Objectifs et moyens des plateformes participatives contemporaines Notre deuxième contribution est un cadre de description des plateformes participatives contemporaines basé sur deux caractéristiques différentes : les objectifs que les organisateurs d’initiatives participatives en ligne cherchent à atteindre et les moyens d’interactions entre les citoyens implémentés. Cette méthode a été proposée à partir d’une série d’entretiens semi-dirigés avec différentes parties prenantes de plateformes participatives. Ces entretiens nous permettent de comprendre les motifs qui poussent des institutions à créer et maintenir des plateformes participatives.

Ensuite, à partir des deux types de différenciations proposés dans ce chapitre, nous proposons des moyens de réduire les problèmes identifiés en nous servant d’outils plus contemporains que ceux implémentés dans les plateformes actuelles. En l’occurrence, nous proposons d’implémenter de l’édition collaborative des contributions et d’utiliser des méthodes issues du Traitement Automatique du Langage (TAL).

Cette contribution a été produite en collaboration avec Rafael Angarita et Valérie Issarny.

Améliorer la participation à grande échelle grâce au TAL La troisième contribution de cette thèse est une méthodologie pour identifier des paires de propositions contradictoires et en accord dans une consultation en utilisant de l’Inférence sur le Langage Naturel (NLI dans le reste de ce résumé). Cette contribution a pour objectifs de faciliter le travail de synthèse des propositions par les administrateurs de plateformes ainsi que l’identification de propositions pertinentes pour un citoyen donné dans le contexte d’une plateforme participative à grande échelle. Ces objectifs sont issus de problèmes précédemment identifiés dans le contexte de la participation citoyenne à la fois en ligne et en personne. Notre méthodologie nous permet de construire des jeux de données permettant d’entraîner des modèles NLI à partir de textes produits directement sur des plateformes participatives, ainsi que d’utiliser ces modèles sur des paires de propositions citoyennes qui sont typiquement plus longues que les phrases communément utilisées pour faire du NLI.

Nous évaluons la pertinence de notre méthodologie pour améliorer la participation citoyenne à grande échelle à partir de deux cas d’usage. Le premier est l’identification de *clusters* de propositions en accord entre elles dans une consultation afin d’automatiser partiellement la synthèse des propositions par les administrateurs de plateformes. Le second cas d’usage est la recommandation, faite aux lecteurs d’une proposition donnée, d’autres propositions contenant des éléments de contradiction afin d’aider les citoyens à identifier les différents points de vue exprimés sur une question donnée.

Cette contribution a été produite en collaboration avec Oana Balalau, Rafael Angarita et Valérie Issarny.

2 Notions préliminaires

2.1 Participation citoyenne

Cette thèse étant orientée vers la participation citoyenne en ligne, il est important de commencer par présenter l’état de l’art sur ce sujet. Nous allons d’abord présenter



Figure 1: Schéma du *workflow* proposé par Parra *et al.* [128] pour décrire les budgets participatifs.

la participation citoyenne en elle-même avant de mettre en avant quelques travaux d'autres chercheurs sur la participation en ligne et sur lesquels cette thèse se base.

La participation citoyenne est un ensemble de pratiques visant à donner aux citoyens un certain degré d'influence sur la gouvernance d'une institution spécifique à laquelle ils sont liés (par exemple, leur ville, leur université, le parlement de leur pays). Dans cette thèse, nous utiliserons la définition de la participation citoyenne donnée par Baum [22], selon laquelle la participation citoyenne est "l'implication des citoyens dans la prise de décision publique". Plusieurs avantages de la participation citoyenne ont été identifiés et sont régulièrement utilisés comme arguments en sa faveur. Un exemple d'avantage est la capacité des citoyens à produire des connaissances originales basées sur leur expertise ou leur expérience. L'inclusion de ces connaissances est considérée comme capitale dans les sociétés contemporaines où les citoyens sont capables de remettre en question ou de compléter l'expertise fournie par les institutions [23, 110]. Différents moyens ont été proposés pour caractériser la participation citoyenne. On peut par exemple citer l'échelle de la participation citoyenne d'Arnstein [15]. Cette échelle est composée de huit degrés de pouvoir citoyen croissant, de la manipulation jusqu'au contrôle citoyen. Chaque degré implique différentes formes d'interaction entre les citoyens et l'institution impliquée, mais aussi entre les citoyens eux-mêmes. Par ailleurs, Parra *et al.* [128] proposent différents espaces de conception pour les éléments informatisés d'un budget participatif à partir d'un *workflow* qu'ils définissent. Ce *workflow* peut être utilisé pour décrire la plupart des initiatives de participation citoyenne. En effet, l'identification des étapes du *workflow* au cours desquelles la contribution des citoyens est demandée dans un processus participatif donné peut être un moyen utile de le décrire. Pour mieux comprendre les différentes étapes potentielles d'un processus participatif, nous reproduisons le schéma représentant ce *workflow* dans la Figure 1.

La participation citoyenne en ligne a fait l'objet de nombreuses études en informatique, en particulier dans les domaines de recherche en TCAO et en IHM, comme l'indique Nelimarkka [117]. Ces études portent sur différents types de contextes participatifs, notamment les systèmes de pétition en ligne [74, 162, 178] ou les budgets participatifs [128, 182]. Les études mentionnées ci-avant sur les pétitions en ligne identifient des corrélations entre le nombre de pétitions sur un sujet donné et le nombre de signatures qu'elles reçoivent, analysent le comportement des utilisateurs les plus actifs à court et à long terme afin d'identifier leurs différences avec les autres signataires de pétitions, et identifient les indicateurs de succès d'une pétition tels qu'ils sont perçus par les contributeurs. Concernant le budget participatif, les au-

teurs identifient les facteurs essentiels incitant les citoyens à rester impliqués dans des campagnes récurrentes après la première année, et analysent les pratiques et technologies utilisées dans les différentes étapes des budgets participatifs en ligne contemporain afin d'identifier leurs failles. De plus, grâce à une revue de la littérature étudiant empiriquement la participation citoyenne, Ianniello *et al.* [75] identifient plusieurs obstacles principaux à la participation et proposent des solutions pour les surmonter. Les trois familles d'obstacles identifiées sont les facteurs contextuels (par exemple, l'attitude des parties prenantes), les dispositions de l'institution (par exemple, le *design* du processus participatif) et la gestion des processus (par exemple, la qualité de la collaboration citoyenne). Certaines conclusions de cette thèse visent à améliorer les moyens offerts par les plateformes participatives pour faciliter la collaboration citoyenne, réduisant ainsi le troisième obstacle identifié.

2.2 Traitement du langage

La recherche en TAL a proposé différentes méthodes qui peuvent être mises en œuvre pour résoudre certains des problèmes de la participation en ligne à grande échelle. Nous définissons ici quelques notions clés liées à l'utilisation du TAL dans cette thèse. Nous présentons également des recherches antérieures en TAL pour la participation citoyenne.

Inférence sur le Langage Naturel : L'Inférence sur le Langage Naturel (NLI) est une tâche qui consiste à prédire, à partir de deux textes, une prémisse et une hypothèse, si : *i*) l'hypothèse peut être déduite de la prémisse (implication) ; *ii*) la négation de l'hypothèse peut être déduite de la prémisse (contradiction) ; *iii*) ou si aucune des deux situations précédentes n'est vraie (neutre). Plusieurs jeux de données annotés ont été proposés pour cette tâche, tels que les corpus Stanford Natural Language Inference (SNLI) [30] et Cross-Lingual NLI (XNLI) [45]. Les systèmes les plus performants sur les *benchmarks* sont des modèles basés sur une architecture "Transformeur", tels que BERT [52] ou RoBERTa [97]. La recherche a montré que le NLI reste un défi complexe [67, 106].

Étude du contenu textuel de plateformes participatives : Pour répondre aux problèmes de surcharge d'informations touchant à la fois les citoyens et les administrateurs de plateformes participatives, plusieurs méthodes basées sur du TAL ont été proposées. Dans [13, 19, 14], les auteurs proposent grâce à différentes méthodes (modélisation thématique, similarité sémantique, etc.) de résumer ou regrouper automatiquement les propositions et les commentaires publiés par les utilisateurs. Dans [149, 150], les auteurs proposent d'extraire des arguments à partir des contributions des citoyens. Ils créent une taxonomie d'arguments et proposent une approche pour identifier les arguments à l'aide de règles syntaxiques. Le module d'extraction d'arguments est ensuite intégré dans un système de recherche par mots-clés.

2.3 Graphes de connaissances

Les graphes de connaissances représentent des faits sous forme de nœuds et de relations [70, 96]. Les deux principaux modèles de graphes sont RDF [103] et les graphes de propriétés [11]. Dans le modèle RDF, les nœuds sont des ressources dotées

d'un identifiant unique sur le web. Les arêtes relient les nœuds en construisant des triplets sujet-prédicat-objet. Les graphes de propriétés, eux, sont des multigraphes étiquetés et dirigés dans lesquels les nœuds et les arêtes peuvent avoir des propriétés sous la forme de paires clé-valeur. Les nœuds y représentent des entités et les arêtes, des relations entre ces entités.

La création d'un graphe de connaissances à partir d'un texte implique l'utilisation de techniques de TAL, notamment la reconnaissance d'entités nommées (NER) [113] et la liaison d'entités nommées (NEL) [154]. Le NER se concentre sur l'identification et la catégorisation des entités dans un texte, telles que les personnes, les organisations et les lieux, entre autres. Le NEL, quant à lui, relie les entités reconnues aux entités d'une base de connaissances. Plusieurs méthodes de construction de bases de connaissances spécifiques à un domaine ont été développées dans divers domaines, dont la mode, la médecine, etc. [21, 141, 140, 161].

3 Étude d'une plateforme participative existante

Les plateformes participatives répondent à divers cas d'usage, par exemple le budget participatif [71, 128], la publication de données ouvertes à des fins civiques [28, 94], la collecte et le traitement des opinions de grands groupes de citoyens [134], ainsi que la consultation et la délibération citoyenne, pour n'en citer que quelques-uns. Dans ce chapitre, nous nous concentrons sur un exemple du quatrième cas : la consultation en ligne "République Numérique"² (abrégé en *RepNum* dans la suite de ce texte). Les consultations en ligne visent à permettre aux citoyens d'exprimer directement leurs opinions sur un sujet particulier – dans notre cas, les différents articles d'un projet de loi – en dehors de tout contexte électoral. Elles sont souvent considérées comme un moyen d'améliorer la participation des citoyens dans un système politique représentatif, bien qu'elles ne soient pas contraignantes pour l'institution organisatrice. Techniquement, les plateformes de consultation en ligne peuvent être considérées comme des plateformes de TCAO permettant aux utilisateurs de modifier de façon collaborative un ensemble de propositions indépendantes en leur ajoutant de nouveaux éléments ou en discutant des éléments ajoutés précédemment. Elles mettent généralement en œuvre au moins les trois formes de contributions suivantes :

1. L'Utilisateur A publie une proposition, qu'elle soit détaillée ou vague, sur la plateforme.
2. L'Utilisateur B commente la proposition précédemment publiée par l'Utilisateur A.
3. L'Utilisateur B vote pour, contre, ou exprime sa neutralité sur la proposition de l'Utilisateur A.

Dans cette partie de la thèse, nous visons à formuler des recommandations pour atténuer les biais dans le *design* de la plateforme afin d'améliorer l'engagement des citoyens dans les plateformes participatives en ligne. Notre objectif est d'aider à construire des plateformes participatives plus équitables sur la base des données produites dans le cadre d'expériences antérieures. Pour ce faire, nous analysons les

²<https://www.republique-numerique.fr/>

contributions et les interactions des utilisateurs sur la plateforme *RepNum*. Cela nous amène à étudier les questions de recherche suivantes, à partir desquelles nous formulons des recommandations pour la conception de plateformes participatives :

QR1: Comment l'attention est-elle distribuée entre les propositions et peut-on identifier des facteurs spécifiques à la plateforme à l'origine de cette distribution ?

QR2: Comment les contributeurs les plus actifs utilisent-ils les différents types de contributions (proposition, modification de proposition, argument, source, vote)?

QR3: Quel est l'impact de la présentation des propositions sur la participation ?

Nous répondons à ces trois questions de recherche de la manière suivante. Tout d'abord, en représentant la consultation sous la forme d'un graphe, nous nous interrogeons sur l'existence de propositions et d'utilisateurs centralisant la conversation à l'aide de l'algorithme PageRank [125] et d'autres métriques. Cette analyse nous permet de comprendre si les contributions ont tendance à être centralisées autour de contributeurs "influents" ou autour de propositions considérées comme intéressantes. Nos résultats nous permettent de proposer des recommandations concernant le *design* de la consultation et l'affichage des propositions dans la plateforme. Ensuite, nous étudions les contributions des contributeurs les plus actifs pour comprendre s'ils ont tendance à se spécialiser dans un type de contribution spécifique (par exemple, écrire des propositions, voter sur les contributions des autres, etc. Nous constatons que les votants les plus actifs ont tendance à être complètement séparés des autres contributeurs actifs qui produisent des contributions textuelles. En outre, nous montrons qu'il existe des contributeurs actifs à la fois spécialisés et polyvalents pour les contributions textuelles. Enfin, nous proposons une analyse statistique des réponses aux propositions en fonction de deux facteurs influençant leur positionnement dans l'interface utilisateur. Ces deux facteurs sont la catégorisation des propositions en fonction de la structure du projet de loi et l'épinglage manuel par les administrateurs de la plateforme de certaines propositions au sommet de chaque catégorie. Nous identifions pour ces deux facteurs des distributions inégales des réponses aux propositions, ce qui nous amène à considérer qu'ils induisent un biais dans la participation. Nous proposons donc des modifications de l'interface utilisateur visant à réduire ces biais.

De plus, afin d'atténuer les biais identifiés dans la conception de cette plateforme participative contemporaine, nous proposons d'améliorer la gamme d'outils mis à la disposition des citoyens et des administrateurs de plateformes pour trouver des propositions spécifiques. Pour ce faire, nous proposons de construire un graphe de connaissances des consultations publiques. Les propositions y sont associées à des sujets identifiés dans leur consultation ainsi qu'à des ressources externes issues du graphe de connaissances DBpedia [18]. Le graphe proposé est générique et peut être utilisé pour la plupart des consultations citoyennes en ligne. Ses contributions sont les suivantes : *i*) l'enrichissement des données d'une consultation publique en ligne avec des ressources liées pour améliorer la contextualisation des textes, faciliter son intégration avec d'autres sources de données liées, y compris d'autres consultations, accroître l'inclusivité, la transparence et la responsabilité, et améliorer la prise de décision ; *ii*) l'analyse des propositions pour déduire leurs sujets afin de mieux organiser, caractériser et comprendre les informations contenues dans les

consultations. En catégorisant les contributions d'une consultation en fonction de leurs sujets, il peut être plus facile d'identifier des tendances et des motifs dans l'opinion publique et de comparer les différents points de vue.

4 Objectifs et moyens des plateformes participatives contemporaines

L'utilisation de plateformes participatives en ligne est devenue plus courante ces dernières années, permettant l'organisation de processus participatifs à plus grande échelle [12, 91], à la fois en ce qui concerne leur nombre de contributeurs et la distance géographique qui sépare ces derniers. Ces plateformes implémentent un large éventail de fonctionnalités pour aider les citoyens à contribuer à un projet participatif. Ces fonctionnalités mettent en œuvre les règles d'un processus participatif, telles que définies par l'institution organisatrice. Elles définissent également le degré d'interaction possible entre les contributeurs d'une plateforme donnée, permettant, par exemple, la création de versions modifiées d'une contribution préexistante. Cependant, Ianniello *et al.* [75] identifient que la mauvaise qualité du *design* d'un processus participatif ou de la collaboration citoyenne a été un obstacle majeur à la réussite de processus participatifs par le passé. Le Chapitre 3 de cette thèse identifie également des défauts dans le *design* d'une consultation en ligne spécifique et de son interface utilisateur, ce qui a entraîné des biais mesurables dans les contributions produites par les citoyens. Par conséquent, nous considérons que les fonctionnalités mises en place sont un facteur clé de la réussite des initiatives participatives en ligne.

Outre les fonctionnalités qu'elles implémentent, les plateformes participatives se distinguent également en fonction de leurs objectifs. L'un des moyens de caractériser les objectifs d'un processus participatif peut être d'utiliser le type de résultat attendu par l'institution. Cette différenciation recoupe en partie celle proposée par l'échelle de participation citoyenne d'Arnstein [15] puisque les degrés les plus hauts de l'échelle – appelés “degrés de pouvoir citoyen” – ne peuvent être atteints si la contribution attendue est simple et individuelle (par exemple, répondre à un sondage). De plus, bien que les administrateurs de plateformes ne soient pas nécessairement conscients des spécificités de chaque degré de participation, ils identifient clairement le type de contributions attendu sur les plateformes qu'ils gèrent et peuvent donc être plus à même de décrire les contributions produites par les contributeurs que le positionnement de leur plateforme sur l'échelle d'Arnstein. Par conséquent, nous considérons que le type de contributions attendu est un critère plus approprié pour différencier les objectifs des plateformes participatives que leur degré de participation citoyenne.

à partir des deux axes de différenciation des plateformes participatives mentionnés ci-dessus – objectifs poursuivis et fonctionnalités implémentées –, ce chapitre répond aux questions de recherche suivantes :

QR1: Quels sont les différents objectifs poursuivis par les institutions responsables de plateformes participatives ?

QR2: Comment les fonctionnalités mises en œuvre dans les plateformes participatives impactent-elles le *design* des processus participatifs ?

Pour répondre à ces questions, nous nous basons sur une série d'entretiens semi-dirigés avec des administrateurs et des concepteurs de plateformes participatives. Les questions que nous avons posé aux participants portaient sur les objectifs que cherchaient à remplir leurs plateformes, les fonctionnalités implémentées et les moyens employés pour récolter et produire une synthèse des contributions. Nous nous basons également sur une étude du fonctionnement de différentes consultations qui se sont basées sur les mêmes plateformes logicielles et qui permettaient la production de contributions textuelles, ainsi que sur les contributions qui y ont été produites. Pour assurer la représentativité de notre panel d'interviewés, nous avons choisi d'interroger des acteurs aux profils professionnels variés, issus de villes de tailles différentes et qui utilisent des logiciels différents pour leurs plateformes participatives. Ces entretiens nous permettent d'identifier quatre objectifs différents et non-exclusifs que les institutions organisatrices de processus participatifs peuvent chercher à remplir : sonder les opinions des citoyens, mettre en place des espaces de dialogue entre les citoyens et une institution, créer des communautés actives de citoyens et leur donner les moyens de porter leur voix dans l'espace public, et co-construire des projets. Nous présentons les caractéristiques de chacun de ces objectifs avant de souligner les éventuels problèmes que rencontrent les institutions à l'origine des plateformes associées. Nous proposons également un second axe de différenciation basé sur les moyens d'interaction implémentés dans les différentes plateformes. Nous distinguons deux types de plateformes. Les premières sont les "boîtes à idées", qui implémentent presque exclusivement des moyens d'interaction entre les citoyens et une institution. Les secondes sont les plateformes "interactives", dans lesquelles de multiples moyens d'interaction entre les citoyens sont implémentés. Nous présentons les possibilités offertes par chacun de ces deux types de plateformes ainsi que les défauts auxquels leurs implémentations contemporaines sont confrontées. À partir des différents objectifs et caractéristiques identifiés lors de nos entretiens, nous proposons des recommandations pour améliorer l'expérience des contributeurs et des administrateurs des plateformes de participation citoyenne, afin de guider les administrateurs et les concepteurs de plateformes dans le développement de futures plateformes participatives. Ces recommandations sont axées sur l'identification des contributions pertinentes grâce à l'apprentissage automatique, l'édition collaborative des contributions, et des fonctionnalités préexistantes qui n'ont pas été mises en œuvre dans le contexte de plateformes participatives.

5 Améliorer la participation à grande échelle grâce au TAL

Dans certains cas, les plateformes participatives permettent aux citoyens de voir et d'interagir avec les contributions d'autres personnes, généralement selon un modèle de collaboration asynchrone. Cette partie de la thèse s'intéresse spécifiquement à ces derniers cas où des interactions de nature collaborative entre les citoyens sont rendues possibles. Cependant, bien que les plateformes participatives permettent de co-construire des projets publics à grande échelle, plusieurs études rapportent que les citoyens et les décideurs rencontrent des problèmes à plusieurs stades du processus participatif. Dans [75], les auteurs identifient plusieurs risques, dont le manque de communication entre l'institution et les citoyens, le manque de transparence des

critères de sélection de propositions, ou la difficulté à atteindre des accords, voire des consensus. Une autre difficulté majeure pour ce type d’initiatives est la surcharge d’information qui peut affecter à la fois les citoyens et les administrateurs de plateformes [38, 132].

Ce travail présente un outil permettant de détecter les équivalences et les contradictions dans les propositions des utilisateurs. Notre outil permet aux utilisateurs de trouver des propositions intéressantes - qu’elles soient équivalentes ou contradictoires - et de se rapprocher de leurs auteurs, tandis que les institutions peuvent s’en servir pour créer des projets composés d’ensembles de propositions *sémantiquement cohérentes*. La cohérence sémantique, par opposition à la cohérence syntaxique mise en œuvre par les algorithmes de gestion de cohérence pour l’édition collaborative (par exemple, [8]), est définie comme l’absence de contradiction entre différents textes. Les contradictions qui nous intéressent peuvent être de différentes natures, allant de la contradiction idéologique à des désaccords numériques [50]. Nous nous concentrons sur des projets composés de propositions émanant de différents utilisateurs. Nous considérons alors qu’un projet est sémantiquement cohérent s’il ne contient aucune paire de propositions contradictoires. Garantir automatiquement la cohérence sémantique aidera les citoyens à construire des projets communs complexes et bien développés en mettant en évidence les contradictions potentielles dans leurs parties constitutives. En outre, la détection de paires de propositions en accord ou contradictoires aidera les responsables de la plateforme à identifier des communautés d’idées et à résumer leurs revendications. Nous présentons une approche basée sur le TAL pour identifier les paires de propositions contradictoires par le biais de la classification de texte. Plus précisément, notre objectif peut être identifié comme une tâche de NLI, également appelée dans la littérature "implication textuelle". Parikh *et al.* [126] définissent la NLI comme "le problème qui consiste à établir des relations d’implication et de contradiction entre une prémisse et une hypothèse". Dans notre cas, nous essayons de prédire si deux propositions sont en accord, contradictoires ou neutres. Nous proposons une méthodologie pour produire des modèles permettant de détecter les propositions contradictoires. Nous appliquons notre méthode à deux consultations citoyennes francophones, *RepNum* et "*Revenu Universel d’Activité*"³ (RUA). Nous avons choisi ces consultations parce qu’elles ont reçu un grand nombre de propositions (respectivement 692 et 1760) et qu’elles ont permis aux contributeurs de soumettre des commentaires pour ou contre une proposition donnée, que nous exploitons pour entraîner les modèles. Par ailleurs, les contributions soumises dans le cadre de ces consultations sont généralement plus longues que les paires de phrases trouvées dans la plupart des corpus de NLI. Par conséquent, nous considérons notre travail comme une étape importante vers l’application du NLI aux textes de grande taille. Pour valider notre approche, nous proposons deux évaluations basées sur des cas d’usage concrets. Tout d’abord, nous évaluons la capacité de nos modèles à identifier des communautés de propositions sémantiquement cohérentes entre elles. Cette utilisation permet d’améliorer la capacité des administrateurs de plateformes à produire des synthèses de consultations à grande échelle grâce à un pré-traitement. Ensuite, réalisons une étude utilisateur comparant la pertinence de nos modèles NLI à la similarité textuelle sémantique, une technique qui a déjà été proposée pour organiser les contributions dans des consultations citoyennes [13, 19], pour une tâche de recommandation de propositions à des citoyens.

³<https://www.consultation-rua.gouv.fr/>

6 Conclusions générales

Le gain en popularité de la participation citoyenne en ligne durant les dernières années a rendu capitale la résolution des problèmes affectant les plateformes participatives afin d'inclure davantage les citoyens dans les prises de décisions publiques. En nous basant sur les problèmes que nous avons identifiés et sur des travaux pré-existants, nous proposons diverses solutions pour améliorer la participation en ligne. Ces solutions sont ont pour objectifs de permettre une plus grande collaboration entre les citoyens et de faciliter la participation à grande échelle.

Les contributions de cette thèse sont les suivantes : *i*) Nous identifions des biais dans le *design* de la plateforme participative associée à la consultation *République Numérique* à partir des interactions entre les citoyens et les contributions. Ces différents biais identifiés nous permettent de proposer cinq recommandations visant à améliorer les plateformes participatives. *ii*) Nous proposons une méthodologie afin de construire un graphe de connaissances pour la participation citoyenne. Nous utilisons ce graphe de connaissances pour permettre une meilleure interopérabilité entre les différentes consultations et pour enrichir les contributions publiées par les citoyens grâce à des ressources externes. *iii*) Nous proposons un cadre de description des plateformes de participation citoyenne basé sur les objectifs que les organisateurs cherchent à atteindre et sur les moyens d'interaction entre citoyens implémentés. Ce cadre de description est basé sur une série d'entretiens semi-dirigés avec différentes parties prenantes de la participation citoyenne en France et en Espagne. *iv*) Nous proposons de faciliter la participation en ligne à grande échelle en identifiant des paires de propositions en accord et en contradiction dans des plateformes participatives. Nous proposons et implémentons une méthodologie visant à produire des classificateurs capables d'effectuer cette tâche sur des contributions citoyennes. Les classificateurs que nous produisons sont utilisés pour évaluer la pertinence de notre méthodologie pour deux cas d'usage différents.

6.1 Perspectives

Les travaux effectués dans cette thèse nous permettent de proposer de nouvelles directions de recherche. Tout d'abord, le graphe de connaissances que nous avons proposé permet déjà une amélioration de la recherche d'informations et de la contextualisation des propositions d'une consultation citoyenne. Cependant, nous identifions les pistes d'amélioration suivantes : une intégration de sources de données externes non représentées sous la forme de graphes de connaissances (par exemple, l'Assemblée Nationale française), et une évaluation empirique par des testeurs indépendants de la qualité de nos annotations reliant les propositions à des ressources DBpedia.

De plus, bien que nous ayons proposé une étude des objectifs des organisateurs de processus participatifs en ligne et des moyens d'interactions qu'ils mettent en place, nos travaux n'incluent pas de citoyens impliqués dans des plateformes participatives. Quelques travaux ont déjà été proposés sur les motivations des citoyens dans certaines plateformes participatives spécifiques [74, 116, 146], mais il n'en existe aucune à notre connaissance s'intéressant à plusieurs plateformes différentes. Par conséquent, nous considérons qu'il serait important que de futurs travaux soient dédiés à l'étude des pratiques et des motivations des citoyens impliqués dans des

processus participatifs en ligne. Pour ce faire, nous considérons que la méthode des entretiens semi-dirigés est la plus adaptée.

Par ailleurs, la participation citoyenne prend des formes diverses et les plateformes participatives ne sont pas les seuls endroits où des institutions recueillent l'opinion de citoyens sur des sujets particuliers. L'exemple le plus fréquent est probablement les réseaux, bien que ces derniers ne soient presque jamais mis en interaction avec les plateformes participatives. Par conséquent, de futurs travaux pourraient être proposés à la fois concernant l'analyse comparative de différents contextes de participation en ligne, et proposant des moyens d'interaction entre ces différentes plateformes.

Enfin, les contributions présentées dans cette thèse sont focalisées sur des aspects spécifiques de la participation citoyenne en ligne, ainsi que sur certains problèmes que rencontrent ses parties prenantes. Cependant, nous n'avons pas encore proposé d'intégration de toutes ces contributions dans une même plateforme participative. Il est donc important de formaliser un nouveau *workflow* pour la participation citoyenne à grande échelle à partir des contributions produites afin de faciliter la collaboration citoyenne et de réduire les problèmes que nous et d'autres chercheurs ont identifiés. Une implémentation sur une plateforme participative *open source* telle que Decidim [51] ou Consul [47] permettrait de fournir un exemple de ce *workflow*.

Contents

1	Introduction	27
1.1	Scope of the thesis	28
1.2	Contributions	29
1.3	Outline of the thesis	31
2	Preliminaries	33
2.1	Definitions	33
2.2	Online citizen participation	35
2.2.1	Citizen participation	35
2.2.2	Analyses of existing online citizen participation contexts	37
2.2.3	Improving online participatory platforms	39
2.3	Large-scale online collaboration	40
2.3.1	Asynchronous editing systems	41
2.3.2	Synchronous editing systems	42
2.4	Knowledge graphs	44
2.5	NLP and machine learning	45
2.6	Conclusion	47
3	Analysis of an existing participatory platform	49
3.1	Introduction	49
3.2	The République Numérique citizen consultation	51
3.3	Does the platform influence how users contribute?	53
3.3.1	Analyzing the distribution of contributions across users	53
3.3.2	Analyzing the factors influencing the contributions	54
3.4	How to mitigate the influence of the most active users?	57
3.5	Does the positioning of <i>Proposals</i> have an impact on participation?	59
3.6	Building an Online Public Consultation Knowledge Graph	60
3.6.1	Topic modeling	60
3.6.2	Entity Recognition and Linking	61
3.6.3	Resulting Knowledge Graph	63
3.7	Threats to validity	63
3.8	Related Work	64
3.9	Conclusion	66
4	Objectives and features of contemporary platforms	69
4.1	Introduction	69
4.2	Methodology	71
4.3	Different objectives of citizen participation platforms	73
4.3.1	Surveying citizens' opinions	73

4.3.2	Dialogue between citizens and an institution	74
4.3.3	Creation of opinion groups or neighborhood councils	75
4.3.4	Co-construction of projects	76
4.4	Different types of platforms for contemporary online citizen participation	77
4.4.1	“Suggestion box”-like platforms	77
4.4.2	Interactive platforms	79
4.5	How to enrich online participatory platforms	80
4.5.1	Collaborative systems	80
4.5.2	Natural Language Processing for large-scale processing of contributions	81
4.6	Threats to validity	83
4.7	Conclusion	83
5	Improving large scale participation using NLI	85
5.1	Introduction	85
5.2	Methodology	87
5.2.1	Problem statement	87
5.2.2	Weakly-labeled data for model training	89
5.2.3	Creating a gold dataset for evaluation	89
5.2.4	Automatically labeling long texts	90
5.3	Datasets	91
5.4	Evaluation	95
5.4.1	Fine-tuned models	95
5.4.2	Insights and error analysis	98
5.5	Assessing the integration of an NLI model into a citizen participation platform	100
5.5.1	Identifying idea communities in a consultation	101
5.5.2	Recommending contributions to citizens during their navigation	103
5.6	Interview with stakeholders	105
5.7	Conclusion	107
6	General conclusions	109
6.1	Contributions	109
6.1.1	Analysis of a participatory platform’s design	109
6.1.2	Building a knowledge graph for citizen consultations	110
6.1.3	Identification of the different objectives and features implemented in participatory platforms	110
6.1.4	NLI for large-scale participation	110
6.2	Future research directions	111
6.2.1	Knowledge graphs	111
6.2.2	Qualitative analysis of interviews with stakeholders of citizen participation	111
6.2.3	Interaction of citizen participation platforms with other participation channels	111
6.2.4	Fine-grained identification of contradictions in proposals	112
6.2.5	Formalized workflow for online citizen participation	112
	Bibliography	113

List of Figures

2.1	CSCW matrix	35
2.2	Workflow for participatory budgeting	36
3.1	Organization of categories in the <i>RepNum</i> consultation	52
3.2	PageRank of users in ascending order.	54
3.3	PageRanks correlation for different reply types.	54
3.4	Correlation of the amount of replies to a proposal for different contribution types.	56
3.5	Correlation of activity for each contribution type.	57
3.6	Steps for building the Online Public Consultation KG.	60
3.7	LDA representation of the topics discussed in a citizen consultation .	61
3.8	Visualization of a generic consultation stored in JSON-LD.	64
4.1	Arnstein’s ladder of citizen participation	70
5.1	A proposal with its arguments in the RepNum platform	92
5.2	Confusion matrix for the RUA dataset using the “XNLI \rightarrow RUA” model	97
5.3	Confusion matrix for the RepNum dataset using the “RepNum + RUA” model	97
5.4	Graph representation of a RepNum section	102

List of Tables

3.1	Possible types of replies for a contribution.	51
3.2	Number of contributions of all types by user profile.	52
3.3	NER and NEL results using DBpedia Spotlight on proposals of the RepNum consultation.	63
3.4	Amount of replies to contributions of all types for each category of the consultation.	67
3.5	Amount of replies to governmental contributions of all types for each category of the consultation.	67
4.1	Questions for the interviews	72
4.2	Profiles of interviewees	72
4.3	Categorization of platforms according to their objectives	73
4.4	Categorization of platforms according to their features	78
4.5	Summary of the issues identified and means proposed to solve them	82
5.1	Distribution of NLI labels for pairs of texts used for models training	91
5.2	Examples of pairs of proposals from the RUA consultation with their gold label	94
5.3	Distribution of NLI labels in the gold dataset	95
5.4	Results of the different models for the RUA gold dataset	96
5.5	Results of the different models for the RepNum gold dataset	98
5.6	Examples of pairs of proposals incorrectly labeled	99
5.7	Distribution of NLI labels for pairs of proposals in the RepNum and RUA consultations	101
5.8	Results of the clustering using the Louvain algorithm	103
5.9	Results of our experiment assessing the suitability of our model for recommendation	104

Chapter 1

Introduction

Contents

1.1	Scope of the thesis	28
1.2	Contributions	29
1.3	Outline of the thesis	31

During the recent decades, citizen participation has become increasingly popular to help citizens take part in political processes. This interest in citizen participation has been the strongest at the local scale (e.g., a city, a neighborhood) [100, 107, 155] but has also a few noteworthy examples at larger (e.g., national) scale [5, 56]. A first milestone concerning the implementation of citizen participation can be found in the participatory budgeting proposed in 1989 by the Porto Alegre municipality [155]. However, the origin of this idea is often traced back to the 1960's with activist groups such as civic rights and women liberation movements [130] in the USA or non-communist left-leaning groups in other countries [109]. Citizen participation regroups an ensemble of practices aiming at introducing some form of citizen input in policy making. These different practices imply multiple types of inputs asked from citizens, from individual responses to a survey to collective assessment of a project. Moreover, the degree of implication and of constraint power given to citizens greatly differ, from descending information of citizens by a governing body to the collective elaboration of a project by citizens [15].

Citizen participation has been a topic of interest for multiple academic research fields, especially in social sciences. Political sciences have proposed multiple analyses on the impact [155] and different interpretations [35] of citizen participation. Communication sciences have studied the discourse which has accompanied different citizen participation initiatives or promotion of citizen participation itself [5, 40]. Sociological studies have analyzed the degree of involvement of different social groups in citizen participation processes [59] and the different reasons motivating their involvement [116, 182]. Multiple other fields such as research in law, education or philosophy have also produced contributions to the study of citizen participation. Finally, the creation of online citizen participation platforms has led to the interest of computer science researchers in different fields. Computer science researchers have studied multiple aspects of online citizen participation, including participatory platforms' Human-Computer Interfaces (HCI) or text mining on citizen contributions.

1.1 Scope of the thesis

Online citizen participation, as well as other topics in the field of civic tech¹, has become increasingly popular in recent years as a tool to better integrate citizens in public decision making. Online citizen participation can take multiple shapes, such as participatory budgeting [71, 128], gathering and releasing civic open data [28, 94] or citizen consultation [83, 148]. Moreover, participatory platforms have addressed a wide range of topics, including public space planning [176], digital transformation of a country [91] or climate change mitigation and adaptation [63, 133]. However, two constants can be identified and are considered here as constitutive of online citizen participation: an online platform where citizens can produce and publish their contributions, and an organizing body² (e.g., municipal authority, university) which initiates the participatory process.

In this Ph.D. thesis, we aim at solving issues faced by both citizens and platforms administrators in contemporary participatory platforms, viewed as social computing systems. Some of these issues have been identified by previous works [38, 75, 132] and others are identified in Chapters 3 and 4 of this thesis. More specifically, we focus on reducing the information overload faced by citizens and platforms administrators and facilitating collaboration among citizens in the context of large-scale online participation. The main reason why we choose to focus mainly on citizen participation at large scale is because we consider the scalability to be one of the main advantage of online participation over its offline counterpart. For example, in 1998, Robert Dahl highlights the fact that letting each citizen express themselves for ten minutes in a city of 10,000 citizens would require more than 200 days at a rate of 8 hours of discussion per day [48]. While this issue is particularly relevant for offline participation, it is largely reduced online thanks to the possibility given to citizens to express or read one another's expression at any time and the absence of the need for a physical space where these contributions to the debate can be expressed. Online platforms therefore allow citizen participation to take place at many different scales, from a neighborhood [146] to multiple countries [20].

To solve the aforementioned issues, we base our works on state-of-the-art contributions from the Human-Computer Interaction (HCI), Computer-Supported Cooperative Work (CSCW) and Natural Language Processing (NLP) fields of computer science research. We propose to answer the two following questions:

1. What are the issues faced by citizens and administrators in online citizen participation platforms?³
2. Based on state-of-the-art computer science research, how can we reduce the issues identified to produce more collaborative and viable at large scale participatory platforms?

Finally, we have to acknowledge the fact that improvements in the technical functioning of participatory platform will not solve by itself the different crises faced by

¹Civic tech refers to the different technologies and related practices aiming at giving a more important role to citizens in democratic processes.

²Often mentioned in this thesis as the "institution"

³We acknowledge that this question has already been tackled by previous research works and aim at further investigating this topic.

representative democracy [62, 99, 163]. Recent demands, for example by the *Gilets Jaunes* [66, 93] or 15-M/*Indignados* [57, 144] movements, for a higher degree of involvement of citizens in public decision making cannot be met only by more effective technical systems for participation but also require fundamental institutional changes to give actual weight to citizens' proposals. In short, this thesis is not written with a techno-solutionist ambition to solve the different problems faced by citizen participation (e.g., lack of representation of minority groups [58], denial of political conflict [112], etc.), but with the objective of detecting existing issues undermining citizen participation in contemporary online platforms and proposing means to solve them.

1.2 Contributions

During this thesis, we produced contributions focused on three different approaches to improve collaboration in online large-scale citizen participation. The said contributions are each the topic of a chapter of this thesis and aim at identifying or solving issues faced by online citizen participation at large scale by leveraging different fields of computer science research.

Analysis of a participatory platform The first contribution of this thesis is an analysis of a specific citizen consultation focused on gathering opinions and new articles for the *République Numérique* law bill. We study the interactions between contributions and contributors of this citizen consultation to identify flaws in the design of the platform. We propose user interface and platform design modification recommendations to mitigate the impact of the flaws identified.

First, using a graph representation of the consultation, we question the existence of proposals and users centralizing the conversation using the PageRank algorithm [125] and other metrics. This analysis allows us to understand whether contributions tend to be centralized around “influential” contributors or around proposals viewed as interesting. Our results allow us to propose recommendations concerning the design of the consultation and the display of proposals in the platform.

Then, we study the contributions made by the most active contributors to understand whether they tend to be specialized in one specific type of contribution (e.g., writing proposals, voting on others' contributions, etc.) or not. We identify a tendency for the most active voters to be completely separate from other active contributors who produce text-based contributions. Moreover, we also highlight the existence of both specialized and versatile active contributors for text-based contributions.

Finally, we propose a statistical analysis of the responses to proposals based on two factors influencing their positioning in the user interface. These two factors are the categorization of the proposals according to the structure of the law bill and the manual pinning by platform administrators of certain proposals at the top of each category. We identify for both factors uneven distributions of replies to proposals, which leads us to consider them as inducing bias in participation. We therefore propose user interface modifications to reduce the said biases.

Additionally, we propose a knowledge graph-based representation of online citizen consultations. This representation allows us to associate proposals with different topics and entities from the DBpedia knowledge base. We apply this representation

to the *République Numérique* online consultation to improve the possibilities to retrieve relevant proposals for a citizen and to summarize contributions for platforms administrators.

This contribution has been published in the following article: William Aboucaya, Rafael Angarita, and Valérie Issarny. “Detecting Obstacles to Collaboration in an Online Participatory Democracy Platform: A Use-case Driven Analysis”. In: 44th International Conference on Software Engineering (ICSE’22). Proceedings of the 2nd International Workshop on Equitable Data and Technology. FairWare ’22. Pittsburgh, Pennsylvania: Association for Computing Machinery, May 2022, pp. 25–33 [1]. The knowledge graph-related part of the chapter has been published in the following article: William Aboucaya, Sonia Guehis, and Rafael Angarita. “Building Online Public Consultation Knowledge Graphs”. In: Text2KG 2023: International Workshop on Knowledge Graph Generation from Text, Co-located with the ESWC 2023. Ed. by CEUR-WS. Hersonissos (Crete), Greece, May 2023, p. 12. url: <https://inria.hal.science/hal-04076743>. [2].

Objectives and features of contemporary participatory platforms Our second contribution is a framework to describe contemporary online citizen participation platforms based on two different characteristics: the objectives pursued by platform maintainers and the features for interaction they implement. This method has been proposed based on a series of semi-structured interviews with participatory platform administrators, designers and technical advisors. These interviews allow us to understand the motivations pursued by institutions when they decide to create and maintain a participatory platform.

First, we propose a differentiation of participatory platforms based on the objectives pursued by institutions organizing participatory processes. The interviews we conducted allow us to identify four different and non-exclusive objectives: surveying citizens’ opinions, organizing dialogue spaces between citizens and an institution, creating citizen groups and giving public recognition to their inputs, and co-construction of projects. For each of these objectives, we present their characteristics and highlight potential issues faced by maintainers of these platforms.

Then, we complete this first differentiation with another one focused on the functionalities for interaction implemented in the platforms. We distinguish two types of platforms. The first one is “suggestion box”-like platforms, which mainly implement possibilities for interaction between citizens and a given institution. The second one is interactive platforms in which multiple features for interaction between contributors are implemented. We present the possibilities given by both types of platforms and the flaw faced by their contemporary implementations.

Finally, based on the two different types of distinction proposed in this contribution, we propose means to reduce the issues identified using state-of-the-art technologies, namely collaborative editing of contributions and Natural Language Processing.

This contribution has been produced in collaboration with Rafael Angarita and Valérie Issarny.

Natural Language Inference for large-scale citizen participation The third contribution of this thesis is a methodology to identify contradictory and entailed pairs of proposals using Natural Language Inference (NLI). This contribution aims

at facilitating the summarization of proposals by platform administrators and the identification of relevant proposals for contributors in the context of large-scale online participatory platforms. These objectives originate from previously identified issues in online and offline citizen participation. Our methodology allows us to build datasets for the training of NLI models using content produced directly on citizen participation platforms and to use these models on pairs of proposals, which are typically longer than one-sentence inputs commonly used in NLI tasks.

After proposing a methodology for identification of entailed and contradictory proposals, we assess it using the contributions produced in two citizen consultations in French language: the *République Numérique* and *Revenu Universel d'Activité* consultations. Our work leverages state-of-the-art transformer architecture-based language models to perform Natural Language Inference on large texts in French.

Finally, we assess the suitability of our methodology to improve large-scale online participation based on two use cases. The first use case is the identification of clusters of entailed proposals in a consultation to partially automate the summarization of proposals by platforms administrators. This use case is assessed using graph-related metrics. The second use case is the recommendation, to the readers of a first proposal, of other proposals which contain elements of contradiction to help citizens identify the different views in a specific debate. This use case is assessed using an empirical evaluation of the quality of proposals recommended.

This contribution has been produced in collaboration with Oana Balalau, Rafael Angarita and Valérie Issarny.

1.3 Outline of the thesis

The remainder of this thesis is structured as follows. Chapter 2 presents the preliminary notions as well as previous work by other researchers used in the rest of this thesis. Chapter 3 is an analysis focused on the interaction of contributors with a specific citizen participation platform. In Chapter 4, we propose a method to describe contemporary online citizen participation platforms based on their objectives and the features they implement. In Chapter 5, we propose and assess a methodology to facilitate large-scale online citizen participation using Natural Language Inference. Finally, Chapter 6 summarizes the contributions of this thesis, proposes some final remarks and presents potential future research directions.

Chapter 2

Preliminaries

Contents

2.1	Definitions	33
2.2	Online citizen participation	35
2.2.1	Citizen participation	35
2.2.2	Analyses of existing online citizen participation contexts .	37
2.2.3	Improving online participatory platforms	39
2.3	Large-scale online collaboration	40
2.3.1	Asynchronous editing systems	41
2.3.2	Synchronous editing systems	42
2.4	Knowledge graphs	44
2.5	NLP and machine learning	45
2.6	Conclusion	47

2.1 Definitions

In this section, we will define the key concepts for the understanding of this thesis. This includes concepts from both citizen participation and computer science, as well as commonly used terms with a specific meaning in the context of this thesis.

Citizen participation: A set of practices aiming at giving citizens a certain degree of influence in the governance of a specific institution to which they are related (e.g., their city, their university, the parliament of their country). In this thesis, we will use the definition of citizen participation given by Baum [22], stating that citizen participation is “the involvement of citizens in public decision making”. This definition does not specify the degree of involvement of citizens, both concerning their decision-making power and the types of contributions they produce. Therefore, we will specify the type of citizen participation we study in each chapter of the thesis as the nature of contributions and the level of empowerment given to citizens may vary significantly between them.

Consultation: Consultation is a specific type of citizen participation in which citizens are asked by an institution to produce contributions on a given topic. Contributions can take various forms, including answers to a survey, textual proposals or videos. The specificity of consultation is the absence of decision power given directly to citizens. Consultations are particularly well-represented in contemporary citizen participation, with common examples being participatory budgeting or neighborhood councils.

Citizen participation platform: An online platform used in the context of a participatory initiative to let citizens produce contributions. These platforms are alternatives to more traditional structures for citizen participation, such as offline neighborhood councils.

Proposal: In the context of a citizen participation platform, a proposal is a text-based contribution which presents an idea, its objectives and potentially details its functioning or expected costs and benefits. Unlike comments, proposals do not reply to another contribution and can be read independently.

Scale: In this thesis, the **scale** of a citizen participation platform refers to the number of citizens involved and the number of contributions they produced. Typically, a large-scale participatory platform would aim at giving voice to the citizens of a large city or a country.

PageRank: The PageRank algorithm [125] is an algorithm aiming at measuring the importance of vertices in a directed graph. In this algorithm, after an initialization of the PageRank of each vertex to a given value, this value is re-computed for any vertex u using the following formula:

$$PR(u) = \sum_{v \in B_u} \frac{PR(v)}{|E^+(v)|}$$

With B_u the set of vertices with at least one edge pointing towards u , and $|E^+(v)|$ the total number of edges outgoing from the vertex v .

The most well-known use of this algorithm is to rank web pages in a search engine based on the number of other pages containing links towards them. However, multiple other use cases have been found for the PageRank algorithm, such as the detection of influential users on Twitter [24].

Collaborative editing: The editing of a single document by several contributors, each of them contributing individually to the final result. Collaborative editing can be performed using various means, from joint writing of a document by two persons in the same room at the same time to online development of an open source software by tens of persons in multiple countries and over a long period of time. The CSCW Matrix (see Figure 2.1) illustrates the diversity of computer-based collaboration, including collaborative editing.

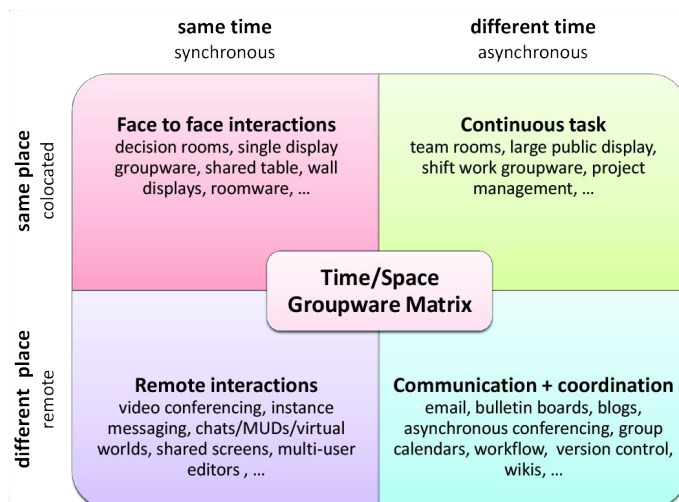


Figure 2.1: Johansen’s CSCW matrix [80] – Image released under the public domain at <https://commons.wikimedia.org/wiki/File:Cscwmatrix.jpg>

Kappa score: Cohen’s kappa score [41] is a measure for the degree of agreement between two persons annotating the same dataset. The kappa score is a number lower or equal to 1, with a score of 1 expressing a perfect agreement between the two annotators, and a score lower than 0 expressing a lower agreement than expected by chance. The score is computed as follows:

$$\kappa = \frac{p_o - p_c}{1 - p_c}$$

With p_o the proportion of items labeled identically by both annotators and p_c the proportion of items expected to be labeled identically by chance.

Token: In NLP vocabulary, a token is a part of a character string. More specifically, in this thesis, tokens are either words, subwords, numbers or punctuation marks. Each sequence of tokens also starts with a specific opening token and ends with another specific closing token.

2.2 Online citizen participation

With the focus of this thesis being on online citizen participation, we first have to present parts of the work done by other researchers on this field. First, we will briefly summarize elements of research on citizen participation itself to define its objectives and identified limits. Then, we will propose a presentation of state of the art in online citizen participation research will focus on two axes: existing analyses of online citizen participation itself or specific platforms, and means to improve contemporary platforms.

2.2.1 Citizen participation

Citizen participation aims at involving citizens in public policy making through a wide range of practices. Indeed, participation has been used with multiple objectives

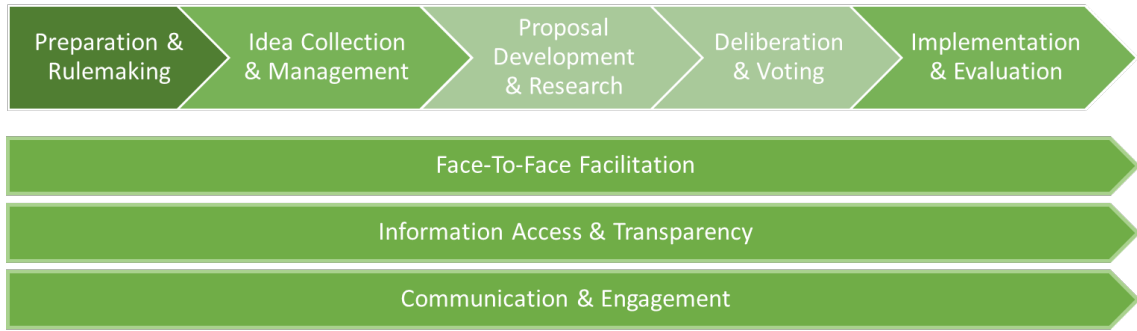


Figure 2.2: Illustration of Parra *et al.*'s [128] workflow for participatory budgeting.

depending on the type of process implemented and the degree of power given to citizens. In [27], Blondiaux briefly compares citizen participation experiences in Porto Alegre, which aimed at redistributing wealth and power to underrepresented citizens through participatory budgeting, and those in England, where the focus was on identifying the preferences of “public services users”. Other benefits of citizen participation have also been identified and are often used as arguments in favor of citizen participation. A first example of such argument is the capability of citizens to produce original knowledge based on their expertise or experience. Inclusion of this knowledge is considered as mandatory in contemporary societies where citizens are capable of questioning or supplementing the expertise provided by institutions [23, 110]. Another benefit of citizen participation is the opportunity for citizens to come together and collaborate in support of common causes, therefore reducing the loss in popularity of collective structures identified since the 1990s [135].

However, multiple limits have also been identified in citizen participation practices. The main issue is a tendency to under-represent minority and vulnerable groups. This issue has been linked to the tendency of deliberative democracy to encourage expressions considered as measured and rational and to disregard more radical and vehement stances [58, 112]. Additional explanations have also been proposed, such as the expected low impact of a participatory process on the final decision or the absence of practical means to encourage the participation of certain groups (e.g., children care services for late meetings) [27]. Finally, in the specific case of online participation, another group is under-represented: computer-illiterate persons, which highlights the need to maintain offline participatory initiatives.

In [128], Parra *et al.* propose different design spaces for computer-based elements in participatory budgeting based on a workflow they define. This workflow can be used to describe most citizen participation initiatives. Indeed, identifying in which steps of the workflow citizen input is requested in a given participatory process can be a helpful mean to describe it. Typically, a participatory process where citizens are given a high degree of control will request citizen inputs in most steps of the workflow, and especially the last ones. To give a better understanding of the different potential steps of a participatory process, we reproduce the schematic representation of this workflow in Figure 2.2.

2.2.2 Analyses of existing online citizen participation contexts

Online citizen participation has been the topic of multiple studies in computer science, especially in the CSCW and HCI fields of research as stated by Nelimarkka [117]. These studies have been made on various types of participatory contexts, including online petitioning systems [74, 162, 178], participatory budgeting [128, 182] or citizen sensing [68, 82]. Studies mentioned above on online petitioning identify non-linear correlations between the number of petitions for a given topic and the number of signatures they receive, analyze the behavior of “power” users in short and long terms to identify their differences with other petition signers, and identify indicators of success of a petition as perceived by contributors. Concerning participatory budgeting, authors identify crucial motivators to keep citizens involved in recurring campaigns after the first year, and analyze the different practices and technologies leveraged in the different steps of contemporary online participatory budgeting to identify their flaws. Finally, research on citizen sensing presents the challenges faced by users of citizen-sourced data, including information overload and misinformation. In addition, it demonstrates the viability of citizen sensing for different topics, such as noise pollution monitoring.

Outside of the field of computer science, multiple works have studied online citizen participation and its different implementations. In her thesis, Laurent [90] proposes an in-depth analysis of the *République Numérique* participatory process from an information and communication sciences perspective. Her works describe the different stakeholders involved in this process and their behavior, from contributors to the platform to media coverage. They also propose an analysis of the design of the participatory platform, examining its significance in the presentation of the debate and exploring the means through which participants can articulate and defend their perspectives [91]. Furthermore, her thesis proposes a thorough analysis of the impact of this participatory initiative on the law-making process and on the final text voted by the French parliament.

The large number of studies focused on online participation has led to the need for means to classify participatory contexts. Kelty *et al.* [84] propose a classification framework for citizen participation based on seven different dimensions. These dimensions are:

Educative dividend: The possibility given to citizens to learn how to participate effectively and to improve their knowledge on the topic of the participatory initiative.

Goals and Tasks: The possibility for citizens to help set the goals.

Resource Control: The capability for citizens to gain control or ownership over the inputs they produce and the means used to produce these inputs. Typically, resource control can be given to citizens in an online participatory initiative through the use of an open source platform and the release of contributions as open data. However, certain participatory initiatives propose other types of practices such as the use of a platform designed with citizens [33, 86].

Exit: The possibility for citizens to disengage from the participatory process at any time and without any penalty.

Voice: The existence of means for citizens influence the outcomes of the participatory initiative and to speak back or criticize the said outcomes in a dialogue between citizens and the institution.

Visible metrics: The existence of relevant links between the inputs produced by citizens and the outcomes of the participatory process.

Affective / Communicative Capacity: The existence of tools for social interaction between citizens (e.g., discussion) and of a sense of belonging to the community of participants from an emotional point of view.

According to their descriptive framework, we consider the participatory platforms studied in depth in this thesis – namely, *République Numérique* and *Revenu Universel d’Activité* – to be strong on the “**Educative Dividend**”, “**Exit**”, “**Visible metrics**” dimensions while being weak on the “**Goals and Tasks**”, “**Voice**” and “**Affective / Communicative Capacity**” dimensions and in-between on the “**Resource Control**” dimension. This distribution is not specific to the two consultations studied here, as we find multiple other examples in online citizen consultations such as Decidim Barcelona¹, Decide Madrid² or New York City’s idea collection³ and participatory budgeting⁴ platforms.

Through a systematic literature review of empirical papers focused on citizen participation, Ianniello *et al.* [75] identify several main obstacles to participation and propose solutions to overcome them. The three families of obstacles identified are contextual factors (e.g., attitude of stakeholders), organizational arrangements (e.g., design of the participatory process), and process management patterns (e.g., quality of the collaboration between citizens). While most of the papers cited in the literature review are focused on offline citizen participation, we consider that many of the conclusions are also relevant for their online counterparts. Related conclusions drawn in this thesis especially aim at improving the means offered by participatory platforms to help citizens collaborate, therefore reducing the third obstacle identified in [75] for citizen participation.

Naranjo-Zolotov *et al.* [116] identify citizens’ motivators in the context of a Portuguese online participatory platforms. Using a questionnaire sent to contributors, they identify that social capital⁵ is a crucial motivator for the engagement of citizens with participatory platforms. They also propose recommendations to help citizens engage with a participatory platform in the long term. Their methodology based on the identification of individual motivators could be used to understand the objectives of citizens in participatory platforms and in particular their motivators with regard to collaboration. Identification of such motivators could be complementary to our work to identify when collaboration is suitable and how it could be implemented in citizen participation contexts.

Glaas *et al.* [63] propose an analysis of contributors’ perception of a specific online participatory platform in Sweden. Their work is based on focus group interviews

¹<https://www.decidim.barcelona/>

²<https://decide.madrid.es/>

³<http://ideas.pbnyc.org/>

⁴<https://www.participate.nyc.gov/>

⁵In the paper cited, the concept of “social capital” is leveraged using Putnam’s definition, i.e., “connections among individuals — social networks and the norms of reciprocity and trustworthiness that arise from them” [135].

with contributors, semi-directed interviews with administrators and a thematic analysis of contributions published on the platform. They identify that the main pitfalls of the platform for contributors are not directly related to its functionalities but rather to the design of the participatory process itself. While our approach differs from theirs – we focus on questions related to the creation of the platforms while they study the experience of contributors –, we consider their findings to be extremely relevant to improve online participatory initiatives independently from their technical basis. While the inclusion of citizens in the design of platforms has already been experimented [33], Glaas *et al.* [63]’s work highlights the importance of leveraging citizens’ experience with the different online participatory platforms to improve contributors’ engagement.

2.2.3 Improving online participatory platforms

Studying the behavior of contributors and their interactions in citizen participation platforms is crucial to identifying potential issues in the design of online participatory initiatives. However, other research has proposed ways to directly improve these platforms based on the issues previously identified. Schroeter [147] proposes to improve citizen participation at the local scale through the use of public interactive devices installed in Brisbane, Australia. The devices displayed ideas proposed by citizens and information on how to submit new proposals using SMS or Twitter, therefore improving the range of possible contexts for citizen participation at local scale. In similar local-scale citizen participation contexts, Yu and Cai [181] suggest enhancing the understanding of citizens and platform administrators by adding geographical annotations to contributions. Kriplean *et al.* [87] propose to modify the types of contributions produced by citizens, asking them to create “pro” or “con” arguments on given issues rather than proposals. Authors identify that such types of human-computer interface tends to allow citizens to form more balanced opinions, acknowledging potential opposing arguments. Cantador *et al.* [34] propose identifying controversial proposals based on different statistical metrics. Their study allows them to identify different controversies types and highlight a lack of interest for certain controversial proposals. However, these papers propose to improve participatory platforms based on elements outside citizens’ textual content. We consider that, in platforms allowing the creation of text-based contributions, analysis of these contributions is particularly useful to reduce information overload and give access to the main points of view and arguments in a debate.

Arana-Catania *et al.* [13] propose utilizing machine learning and NLP to improve citizen participation systems. Their exploratory study focuses on the relevance of machine learning to improve citizen participation platforms based on four different use cases. These use cases are the summarization of comments using text generation, the clustering of users using their topics of interest, and the identification of similar proposals using automated tagging or finding the most representative proposals for a given topic. NLP analysis of contributions can be an interesting mean to help reduce the information overload identified in participatory platforms [38, 132] and to foster citizen collaboration at large scale. Indeed, their user-centric metrics for assessment of their use cases (e.g., time to perform a given action with or without their tool) show that the implementation of their tools on the “Decide Madrid” platform drastically facilitate the retrieval of information for the different proposed

scenarios. We consider that the tools they propose could help citizens regroup to create collective contributions based on the similarity of their proposals or user clustering.

Athanasopoulos [16] introduces a framework for online direct democracy with a focus on privacy and scalability. Their framework allows the organization of votes on discussion topics and is based on the representation of participants as vertices in a distributed social network. The proposed network architecture allows both achieving better scalability in the number of participants than a centralized one, and avoiding the existence of a central authority owning the data. This framework could be implemented in online participatory platforms to ensure the privacy of participants and scalability up to the country's population. However, it is necessary to extend the means of participation offered to a wider scope than the sole (dis)approval of proposals by implementing features supporting the other phase of the generic workflow presented in [128].

The development of online citizen participation outside of dedicated platforms has been the subject of prior studies. Angarita *et al.* [9] propose a middleware to help citizens interact with one another and with other relevant entities (e.g., connected objects, institutions, etc.) through any social network and to identify and solve public issues. Saldivar *et al.* [143] analyze the behavior of contributors on a city-scale online participatory initiative where citizens were allowed to contribute both using a dedicated platform and a Facebook group, with the content posted on one of these platforms being mirrored on the other. However, their results show that even though the existence of this Facebook-based version of the consultation increased the number of users, a vast majority of the contributions were still made using the dedicated platform. Therefore, we consider that dedicated platforms are still the most relevant place to study online citizen participation and that their technical improvement is a key factor to involve citizens in public decision making more efficiently. Moreover, we stress a major advantage of platforms dedicated to citizen participation, especially those with open-source software, which is the possibility for the organizing institution to implement new features for contributors or administrators (e.g., collaborative editing of contributions) independently from a third-party actor – in Saldivar *et al.*'s case [143], Facebook.

2.3 Large-scale online collaboration

The concept of online collaboration at large scale is largely used in this thesis. Moreover, our findings in Chapter 4 lead us to consider online collaborative editing (CE) as a crucial mean to let citizens produce common contributions. Indeed, (computer-supported) CE systems allow people to work (online) on a common task, and therefore simplify the creation of projects of all sizes. Online collaboration also allows contributors from almost everywhere to participate without having to be physically present at a given place, neither at a given time. Therefore, we propose a summary of the two types of collaborative editing we differentiate in this thesis: asynchronous and synchronous editing.

2.3.1 Asynchronous editing systems

In an asynchronous CE system, updates are explicitly shared by their creator and received by other participants upon request. This approach allows users to perform major updates to a document without compromising the other contributors' work. However, this challenges the automation of conflict resolution, which is generally handled by the users so that the document remains coherent with all the conflicting versions. This then requires a quite high level of computer literacy from the contributors who need to understand the issue of conflicting updates.

CE systems have taken various shapes over time, from sharing document updates via emails to massive online collaboration platforms, like GitHub or Wikipedia. They have also been used in a wide range of situations, both in software development and in natural language document writing.

In the context of software development, the most common tools for asynchronous editing are Version Control Systems (VCSs), which are defined in [37] as “*a system that records changes to a file or set of files over time so that [one] can recall specific versions later*”. They are frequently preferred to synchronous file editors for various reasons:

- It allows users to gain full control over their replica of the code they are working on. Assuming that the code runs without issues on the local machine when an update is received, this ensures that any bug or error encountered during development has been caused by subsequent changes made by the user and, therefore, simplifies the debugging process.
- Since updates in software development are typically bug fixes, new features, or improvement of existing ones, adding a single character or line is often not significant enough to be shared with other contributors, and may even create unpredictable behaviors and errors. The possibility to create updates, and label and share them when users want, gives contributors a higher control over the process.
- In the specific case of open source software development, habits and customs are shown to have a strong impact in certain communities [61, 152] and thereby influence the choice of the hosting platform for a software development project. Furthermore, the fact that the asynchronous GitHub, one of the biggest online storage platforms for code repositories, claims that 20.5 million new users have created an account in 2022 [121], suggests that choosing GitHub potentially gives greater visibility to the project.

Asynchronous editing systems for software development can be suitable tools for citizen participation since they allow the creation of *branches* (i.e., alternative versions) within a single project. This feature, when applied to participation, allows citizens to collectively create several proposals –at different stages of their development– for a same project, to fully develop them and then to choose the most suitable alternative.

Asynchronous CE systems also have a strong presence for the CE of documents in natural language, mainly in wiki engines⁶. Indeed, even though the concept of wiki

⁶Collaborative software running a wiki platform

is not intrinsically linked to asynchronous editing, it is clear that the most popular wiki platforms use this approach. This tendency can be explained by the fact that the MediaWiki engine, which powers certain of the most popular wiki platforms implements an asynchronous protocol for sharing updates. The MediaWiki engine, which is used by platforms such as Wikipedia, FANDOM or wikiHow, allows users to edit the content of an article on a page separate from its main view and then send the changes to the server as a revision, with a description of the proposed changes. A change history allows tracking of the different revisions of a document and to possibly restore previous versions or remove the impact of an intermediary one. The efficiency of wikis as a tool for collaboration and knowledge aggregation has been studied in various contexts, like school projects [115, 158], teaching [119, 127, 169] or corporate knowledge management [72]. However, the asynchronous approach has sometimes been seen as an obstacle to collaboration in certain contexts, like collaborative lesson development among pre-service teachers [124].

Overall, asynchronous editing systems tend to be the most appropriate choice when the object of the collaboration requires one of the following functions:

- **Need for explicit versioning:** In many situations, the existence of labeled or fully described versions of a project is mandatory –e.g., in order to trace the history of changes of a paragraph or a software feature. In these cases, choosing asynchronous editing can be an advantage since these systems require contributors to explicitly mark the end of a contribution, and therefore simplify their labeling.
- **Possibility to execute or compile the edited documents:** When the edited document contains computer code, which is meant to be interpreted or compiled, it is often mandatory to adhere to its syntax and semantics to avoid errors. Consequently, using an asynchronous CE system, the developer gains full control on the version of the project they work on, therefore avoiding the risk of compiling it while someone else is making changes which may create syntax and/or semantic errors.
- **Public document edition:** In cases where write access to a document is not restricted to a specific group of users (for example, most wiki websites allow everybody to edit their pages), asynchronous editing may be the only viable option. Indeed, in such a case, there is a non-negligible risk of vandalism [3, 151] and of edit wars [159, 180]. These problems are likely easier to solve if the contributions are large updates rather than atomic changes, since a detailed version history simplifies the recovery of a former undamaged version. However, certain algorithms for synchronous editing that allow the identification and reversion of long updates [160, 173] can also be considered as a solution for such issues.

2.3.2 Synchronous editing systems

To the best of our knowledge, synchronous CE has been implemented for the first time at the Augmented Human Intellect (AHI) Research Center of Stanford Research Institute, in 1968 [55]. Changes made using a synchronous CE system are sent to the participants as soon as they are done or after a short delay. This method allows the

users to see each others updates as they are issued, which facilitates collaboration. Moreover, since these systems are meant to handle simultaneous work on the same document, they are able to resolve writing conflicts automatically, which lowers the level of computer literacy required to use such systems. However, conflict resolution requires specific algorithms designed to maintain consistency among the different replicas of the document edited by each user. These algorithms are grouped into two categories: (1) *Operational Transformation* (OT)-based algorithms [54] and (2) *Conflict-free Replicated Data Types* (CRDT)-based algorithms [153]. We will not detail the technical differences between these two categories here since, in this thesis, synchronous editing systems can be considered as black boxes. We will simply outline the main characteristics of these two categories to help differentiate them.

OT is a class of algorithms designed to maintain consistency over multiple nodes handling replicas of the same object (a database, a document, an array, etc.). It is based on the control and transformation of concurrent operations to achieve the eventual consistency of replicas. It has been designed to solve the potential inconsistency produced by two operations that can not be causally ordered. In order to detect whether two operations are concurrent or not, several methods can be implemented; the most frequent of them is the use of a server assigning timestamps to operations, allowing sites to reconstruct the history by themselves. In an OT-based system, the positions at which received operations should be applied are transformed on each site, based on concurrent local operations. However, most of these algorithms have been proven incorrect during the past years [76] since they do not achieve convergence in certain cases.

As an alternative to OT, CRDT has been defined in [153] as a data structure allowing strong eventual consistency between nodes of the network without any coordination mechanism. Systems that implement this structure send data with an overhead that enables every node in the network to place elements (e.g., characters of a text) at the same position. CRDTs do not need a central server by design and are therefore more appropriate in contexts where there is no trusted entity able to store and process contributions. This property comes from the commutativity of operations, or of the merging function in the case of state-based CRDTs, allowing them to be executed independently of the context with the same result. However, some of the proposed optimizations may need a centralized architecture.

Both of the aforementioned coherence algorithm classes have specific implementations for synchronous CE. However, the proven correctness of CRDTs is a clear argument in their favor.

Nowadays, synchronous editors are mostly used for editing documents in natural language through editors such as Google Docs, Etherpad or MUTE [118]. These editors allow users to create and edit documents like they would do using a regular text processor plus the specific features proposed by synchronous CE. As shown in [123], these editors give users new means of collaboration using recognizable patterns of work. However, the direct sharing of contributions to a document can also be the source of inconveniences for the users, from the complexity induced by collaboration rules of a group to the difficulty to share contributions when they are yet not finished and do not fully represent the ideas their authors want to share [170]. The latter can be a major obstacle for synchronous collaboration in citizen participation circumstances, especially if contributors are supposed to share ideas or knowledge.

Other synchronous CE systems exist, such as synchronous wiki engines (Nuclino⁷, Slite⁸, Confluence⁹, etc.). These systems, like their asynchronous counterparts, allow users to modify the content of a web page directly through its interface. Their functioning is very similar to that of document editors mentioned above. Their main advantage is to avoid asking users to do conflict resolution since it is done automatically. However, they do not allow the creation of a semantically labeled change history.

Finally, we can mention the fact that synchronous editing systems have been proposed for software development [65, 69]. They generally allow users to see each other's code as they are writing it but it is not run by others' computers as long as it is not syntactically correct or marked as ready to use by the author. However, these systems are nowadays rarely used, abandoned or still in prototype stage.

Summarizing, the main situations where synchronous collaborative editors tend to be preferred to their asynchronous alternatives are:

- **Need for a solution that is easy to use:** In cases where the contributors are not supposed to have a high degree of computer literacy, reducing the difficulty of use is of paramount importance. Therefore, automated sharing of contributions and conflict resolution are important features in this kind of situation. Moreover, using a synchronous editor guarantees to the user that their contributions are applied to the most recent version of the project.
- **Necessity for contributors to be able to do joint writing:** In certain situations –e.g., when a user is supposed to correct spelling mistakes or when a text is supposed to reflect two persons' perspectives on a subject–, synchronous editing is the most viable option available. Apart from the possibility to do joint writing, synchronous editing also gives contributors the possibility to send feedback to a co-author as they are writing.

2.4 Knowledge graphs

Knowledge Graphs (KGs) represent facts in the form of nodes and relationships [70, 96]. They use a graph-based data model to capture knowledge in scenarios that involve integrating, managing and extracting data from diverse sources [120]. The two main graph models are RDF [103] and Property Graphs [11]. RDF is a W3C standard RDF for data interchange on the Web. Nodes are resources with a unique identifier on the Web. Edges link the nodes by building subject–predicate–object expressions. Property graphs are directed labeled multigraphs in which nodes and edges can have properties in the form of key-value pairs. Nodes represent entities and edges represent relationships between those entities. Both RDF and Property Graphs have their advantages and disadvantages, depending on the use case [6, 53].

The creation of a KG from text involves the use of Natural Language Processing (NLP) techniques, including Named-Entity Recognition (NER) [113] and Named-Entity Linking (NEL) [154]. NER focuses on identifying and categorizing entities within text, such as people, organizations, and locations, among others. NEL, on

⁷<https://www.nuclino.com/>

⁸<https://slite.com/>

⁹<https://www.atlassian.com/software/confluence>

the other hand, connects the recognized entities to entities within a knowledge base. Several methods for constructing domain-specific KGs have been developed across a range of domains, including fashion, medicine, and more [21, 141, 140, 161].

Barroca *et al.* [21] applied Named Entity Recognition (NER) and Named Entity Linking (NEL) techniques on product textual descriptions to enhance a fashion-related KG. The process involved identifying entities, such as materials used in the product, within the descriptions and linking them to their corresponding, unique nodes in the KG. Rincon-Yanez *et al.* [141] proposed a methodology for creating KGs related to novels. This methodology involves using NLP techniques such as NER and NEL to generate simple triples as a starting point, and then enhancing them using external open knowledge sources, such as DBpedia, to build a comprehensive KG. This methodology was tested on a novel, resulting in the extraction of characters and their relationships. In a similar vein, Alpizar-Chacon and Sosnovsky [7] proposed a solution which automatically labels book's index item in relation to the main subject of a textbook area. These index concepts can be associated with the central themes of the book, common concepts, topics related to but distinct from the book's subject area, or unrelated to the book's subject area altogether.

Reklos and Meroño-Peñuel [140] developed an approach for identifying causal relationships in medical publications. Their approach involves detecting causal sentences, extracting entities, and using this information to build a KG that represents the causal relationships. For example, a node representing the cause of an illness would be connected to a node representing the illness.

Another relevant application of KGs is the analysis of parliamentary debates, as demonstrated by Tamper *et al.* [161]. They proposed a method for constructing and enhancing KGs of plenary debates in the Parliament of Finland. The methodology involves preprocessing the text, including lemmatization to normalize the text and subject indexing to describe the text succinctly, due to the particularities of the Finnish language. NER and NEL techniques are then applied to extract named entities and link them to the KG.

We consider knowledge graphs to be a great mean to improve online citizen participation. For example, representing online participatory initiatives as knowledge bases would allow an enrichment of the contributions using preexisting external resources.

2.5 NLP and machine learning

The NLP field of research has implemented different methods which can be leveraged to improve certain issues faced by large-scale online citizen participation. In this section, we define certain key notions related to the use of NLP in this thesis, and particularly in Chapter 5. We will also present previous works on NLP and machine learning research used as background work for this thesis.

Textual Entailment Recognition / Natural Language Inference: Textual Entailment Recognition (TER), also known in the literature as Natural Language Inference (NLI), is the task of deciding, given two texts, a premise and a hypothesis, if: *i*) the hypothesis can be inferred from the premise (entailment); *ii*) or if the negation of the hypothesis can be inferred from the premise (contradiction); *iii*) or if neither of the two situations is true (neutral).

Several annotated datasets have been proposed for the task, such as the Stanford Natural Language Inference (SNLI) Corpus [30] and the Cross-Lingual NLI Corpus (XNLI) [45]. We note that on the SNLI task, inter-annotator agreement varied between moderate and substantial. This is a consequence of the task requiring a common understanding of language and background knowledge. Moreover, the exact definition of what can be inferred from a piece of text is left to the understanding of each annotator. The GLUE [168] and SuperGLUE [167] benchmarks, two standard benchmarks for general-purpose language understanding systems, include textual entailment as a challenge. The best performing systems on the benchmarks are Transformer-based models, such as BERT [52], RoBERTa [97] or PaLM [39]. Research has shown that the performance of classifiers on such datasets is misleading and textual entailment remains challenging [67, 106]. In many cases, the classifiers learn simple heuristics to predict a class, such as token overlap between entailed sentences or the presence of negation for contradictory sentences.

Analyzing the textual content of participatory platforms. Participatory platforms are a meaningful way of engaging citizens and finding innovative solutions to societal struggles. One of the success criteria for collaborative platforms is the number of participating citizens. However, decision-makers, such as governments, often do not have the means to deal with high volume and complex citizen inputs [38]. In addition, citizens might struggle to assess if their contributions bring new elements to the conversation and might be overwhelmed by the other contributions [132]. Several techniques have been proposed in the literature to alleviate this problem. In [13], the authors propose using topic modeling to generate tags for each proposal to facilitate the search by keywords, extractive summarization for the comments associated with proposals, and user clustering based on the contributions. In [19], proposals are clustered using the Word Moving Distance, while in [14] comments are summarized using abstractive summarization models. In [149, 150], the authors propose extracting arguments from user input. In [150], the authors create a taxonomy of arguments and propose a rule-based approach using syntactic patterns to extract arguments. In a follow-up work [149], the argument extraction module is integrated into a keyword search system. Different from previous works, in Chapter 5 of this thesis, we focus on a more fine-grained understanding of user input by retrieving entailed and contradictory proposals, which should be grouped or reconciled. In addition, we use state-of-the-art techniques for natural language processing, and we evaluate our approach on two major national citizen consultations.

Semantic Textual Similarity. Semantic Textual Similarity (STS) measures the degree of semantic equivalence between sentences. This task is fundamental in NLP and has essential applications in measuring the success of other tasks such as machine translation, summarization, or question answering. Semantic textual similarity is related to textual entailment and paraphrase detection but differs from these two tasks in that it does not perform a binary classification but captures the degree of similarity. Several supporting datasets have been proposed [36, 104, 111]. Similarity is usually measured on a scale of 0 to 5. Each gradation is explained to human annotators: completely dissimilar sentences are given 0; equivalent sentences are given 5; and, for example, sentences that are roughly equivalent but have important

different or missing details are given 3. There is no clear guideline for dealing with contradictory sentences, which may well receive high similarity scores [36]. One of the most successful models for computing textual similarity are siamese Transformer networks [137].

2.6 Conclusion

This chapter lays the foundation for the rest of my Ph.D. thesis by introducing key concepts and providing background information relevant to these concepts. The first section highlights certain key definitions used in all or parts of this document. We then propose a presentation of preexisting analyses of online citizen participation from the HCI and CSCW fields of research, as well as a summary of the means proposed to solve already identified issues in participatory platforms based on state-of-the-art technologies. We also present key notions to understand online collaboration with a focus on the synchronicity of collaborative systems. Then, we summarize a series of concepts and preexisting works from knowledge graphs research. Finally, we conclude this chapter with a presentation of NLP and machine learning concepts and background works used in the following parts of this thesis.

The following chapter of this thesis will be focused on the presentation of the first contribution of my Ph.D.: an analysis of a specific citizen consultation platform based on the interactions between its contributors and their contributions.

Chapter 3

Analysis of an existing participatory platform

Contents

3.1	Introduction	49
3.2	The République Numérique citizen consultation	51
3.3	Does the platform influence how users contribute?	53
3.3.1	Analyzing the distribution of contributions across users	53
3.3.2	Analyzing the factors influencing the contributions	54
3.4	How to mitigate the influence of the most active users?	57
3.5	Does the positioning of <i>Proposals</i> have an impact on participation?	59
3.6	Building an Online Public Consultation Knowledge Graph	60
3.6.1	Topic modeling	60
3.6.2	Entity Recognition and Linking	61
3.6.3	Resulting Knowledge Graph	63
3.7	Threats to validity	63
3.8	Related Work	64
3.9	Conclusion	66

The content of this chapter has been published in the FairWare’22 workshop proceedings [1] under the name “Detecting Obstacles to Collaboration in an Online Participatory Democracy Platform: A Use-case Driven Analysis”, with the exception of Section 3.6. This section is based on the results of our paper “Building Online Public Consultation Knowledge Graphs” [2], written with Sonia Guehis and Rafael Angarita and published in the Text2KG’23 workshop proceedings.

3.1 Introduction

Participatory platforms address various applications, e.g., participatory budgeting [71, 128], releasing open data for civic purposes [28, 94], gathering and processing

the opinions of large groups of citizens [134], and citizen consultation and deliberation, to name a few. In this chapter, we focus on an example of the fourth case: the “*République Numérique*” (Digital Republic, *RepNum* for short in the following) online consultation (§ 3.2). Online consultations aim at allowing citizens to directly express their opinions on a particular topic—in our use case, the different articles of a bill on France’s digital transformation—outside of any electoral context. They are frequently considered as a way to improve the participation of citizens in a representative political system although they are non-binding for the organizing institution. Technically, online consultation platforms can be seen as CSCW platforms allowing users to collaboratively edit a set of independent proposals by adding new elements or by discussing previously added ones. They generally implement at least the three following forms of contributions:

1. User A submits a proposal, detailed or vague, to the platform.
2. User B comments on a proposal formerly submitted by User A.
3. User B votes for, against, or neutrally, on A’s proposal.

The *RepNum* platform has already been the subject of several academic papers, notably in the French-speaking social science community. Existing analyses focus mainly on the discourse accompanying the implementation of the platform [5] and on the interaction possibilities offered to citizens by the platform design [91]. However, to the best of our knowledge, there does not exist any publication focused on the technical analysis of the impact of platform design on the way users contributed to the *RepNum* consultation. There exist technical analyses of the user contributions to similar platforms such as Decidim [12] or more general collaborative systems like Google Docs [123]. These studies focus on the inference of collaboration models from the users’ interactions. Instead, we aim at eliciting recommendations mitigating bias in the platform’s design to enhance the engagement of citizens in online participatory platforms. Our objective is to help build more equitable participatory platforms based on the data produced by previous experiences. To do so, we analyze the users’ contributions and interactions in the *RepNum* platform. This specifically leads us to investigate the following Research Questions out of which we draw recommendations for the design of participatory platforms:

RQ1: How is the attention distributed among proposals and are platform-related factors at the origin of this distribution? (§ 3.3)

RQ2: How do the most active contributors use the different types of contributions (proposal, modification, argument, source, vote) and how can we mitigate their influence on the consultation? (§ 3.4)

RQ3: What is the impact of the presentation of proposals on the participation? (§ 3.5)

Moreover, to mitigate biases identified in this contemporary participatory platform’s design, we propose to improve the range of tools available to citizens and platforms administrators to find specific proposals. To do so, we propose to build a knowledge graph of public consultations in § 3.6. The graph proposed is generic

Table 3.1: Possible types of replies for a contribution.

Initial contribution \ Can be replied with...	Argument	Modification	Source	Vote
Argument				×
Modification	×		×	×
Proposal	×	×	×	×
Source				×

and can be used on most online citizen consultations. Its contributions are the following: *i*) the enrichment of an online public consultation data with linked data to improve the contextualization of texts, facilitate its integration with other linked data sources including other consultations, increase transparency and accountability, and enhance decision making; *ii*) the analysis of proposals to infer their topics for better organization and understanding of the information contained within the consultations. By categorizing consultations based on topics, it can be easier to identify trends and patterns in public opinion and to compare and contrast different viewpoints.

All code used in the production of the analyses is available at <https://github.com/WilliamAboucaya/repnumbehavioranalysis> for Sections 3.3, 3.4 and 3.5, and at <https://github.com/WilliamAboucaya/repnum-kg> for Section 3.6.

We then analyze the threats to validity (§ 3.7) and position our contributions with respect to related work (§ 3.8). A summary of our contributions and perspectives for future work conclude the chapter.

3.2 The République Numérique citizen consultation

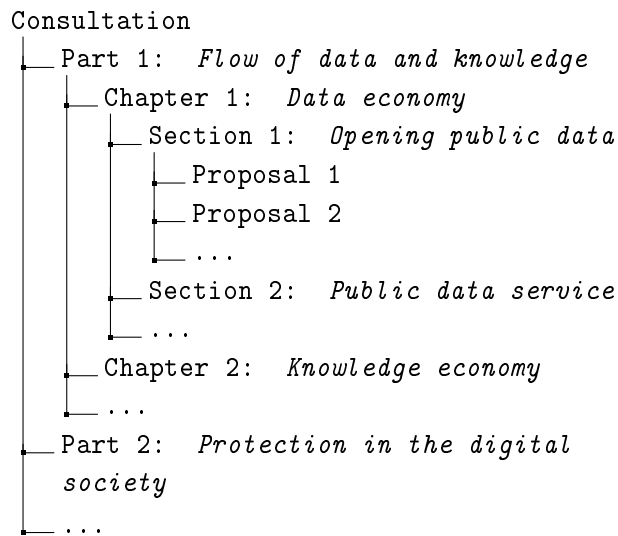
The consultation for the *RepNum* bill was held from 09/26 to 10/18 2015 and primarily took place through the dedicated online platform [43]. The main objective of the (essentially online) consultation was to allow contributors to propose amendments –called *Proposals*– to the bill before its examination by the French Parliament, while giving no guarantee that the amendments would be integrated.

The *RepNum* platform builds upon the proprietary solution for online consultations of the Cap Collectif company¹. Using the *RepNum* platform, contributors register using their real name or a pseudonym. They also specify their *profile*, that is, whether they are *citizens*, *institutions*, *non-profit* or *for-profit* organizations. Over the last few years, other consultations have taken place in France, particularly concerning vaccination [95, 171], and in other countries such as Australia [32] and Canada [98]. However, we chose to analyze the *RepNum* consultation because it received a large number of contributions, and the platform allows more diverse types of contributions (see Table 3.1) compared to its main alternatives, e.g., CONSUL [47] or Decidim [51]. Distinguishing features include the addition of a source (e.g., a newspaper article) to a preexisting proposal, and the amendment of a proposal through a modified version. We argue that such functions are vital in a collaborative

¹<https://www.cap-collectif.com/>

Table 3.2: Number of contributions of all types by user profile.

PROFILE	USERS	PROPOSALS	MODIFICATIONS	SOURCES	VOTES	ARGUMENTS
Citizen	13,492 62.99%	400 57.803%	950 68.493%	369 88.916%	112,322 76.080%	4,627 77.245%
Non-profit organization	393 1.831%	110 15.896%	157 11.319%	14 3.373%	1,955 1.324%	291 4.98%
Institution	199 0.927%	22 3.179%	47 3.389%	11 2.651%	1,141 0.773%	62 1.035%
For-profit organization	154 0.717%	13 1.879%	27 1.947%	4 0.964%	1,043 0.706%	66 1.102%
Government	1 0.005%	30 4.335%	0 0%	0 0%	0 0%	0 0%
Others	7,225 33.661%	117 16.908%	206 14.92%	17 4.096%	31,176 21.117%	944 15.760%
Total	21,464	692	1,387	415	147,637	5,990

Figure 3.1: Organization of proposal categories with examples of categories from the *RepNum* bill (translated from French).

and confrontational situation to help solve potential disagreements and help everyone form an informed opinion on a particular topic, as described in [101]. We further highlight that the functions offered by the *RepNum* platform match the ones of the generic interaction workflow supporting participatory budgeting presented in [128] (see Figure 2.2).

Statistics indicate that the consultation has reached a diverse audience larger than only experts in IT and companies of the field [42, 77, 78]. However, the way users have contributed varies greatly depending on their profile. Table 3.2 illustrates how the different types of contributions are distributed among the different user profiles. For example, it emphasizes that certain profiles, particularly non-profit organizations and the French government, have published many more *Proposals* than others per individual, especially citizens.

All the contributions to this consultation and related metadata are available online as open data at <https://www.data.gouv.fr/fr/datasets/consultation-sur-le-projet-de-loi-republique-numerique>. This enables us to conduct an in-depth analysis of how users engage based on their profile and types of contributions, from which we are able to derive platform pitfalls and recommendations for the platform design to enhance the user engagement. *Contributions* relate to one

of the bill *categories* (i.e., the specific part, chapter, and section, according to the overall structure of the bill depicted in Figure 3.1). As detailed in Table 3.2, the dataset gathers 156,121 contributions from 21,464 unique authors. However, the dataset omits a key feature of the platform: it pins the *Proposals* from the government –i.e., part of the initial law bill– on top of every section of the consultation, while the government profile is of type "*institution*" in the dataset. Thus, to distinguish the government *Proposals*, we pre-processed the dataset to replace their profile "*institution*" by "*government*".

3.3 Does the platform influence how users contribute?

3.3.1 Analyzing the distribution of contributions across users

In order to identify which contributors have received the most attention, we apply the PageRank algorithm [125] that measures the importance of vertices in a directed graph through the number of edges pointing towards them. Various papers introduce PageRank adaptations to detect influential users in social networks, while accounting for the specifics of the underlying platform [24, 174]. We specifically use the weighted PageRank algorithm [179] that fits well the interactions implemented in the *RepNum* platform. We built four weighted directed graphs representing the interactions between contributors using the four types of replies (i.e., *Argument*, *Source*, *Modification*, *Vote*, interactions between these contributions are presented in Table 3.1) to a contribution –*Proposals* are ignored since they cannot follow from another contribution. We distinguish the different types of contributions to avoid giving the same weight to two different contributions that do not have the same purpose and effort-requirement (e.g., a 1-click vote *vs* a detailed argument). It also allows us to analyze if the PageRanks of a user correlate for different types of contributions (see § 3.3.2). Hence, each of the four graphs focuses on a single type of contribution where: the vertices are the users/contributors to the platform, and a directed edge from a contributor *A* to a contributor *B* represents a reply contribution of the chosen type from *A* to a contribution from *B*. Each graph is defined as $G(V, E)$ with V (resp. E) denoting the set of vertices (resp. edges). We note $E^+(v)$ the set of edges outgoing from a vertex v . Then, $w : E \rightarrow \mathbb{R}$ returns the weight of a given edge where the weight of an edge e outgoing from a vertex v is computed as:

$$\forall v \in V, e \in E^+(v) : w(e) = \frac{1}{|E^+(v)|}$$

After applying the PageRank algorithm to the four graphs, we filtered each one of them by removing the users who did not publish any contribution that can generate replies of a given type from the associated graph (e.g., a user who did only publish *Sources* cannot receive an *Argument*). We then ordered the PageRanks in ascending order (see Figure 3.2). For each graph, we can clearly identify a small minority of contributors whose activities generate the most replies, as represented by an exponential growth of the PageRank. We argue that the existence of this small minority whose *Proposals* centralize the vast majority of replies of a specific type tends to alter the nature of the consultation for most users. Indeed, since most users receive little to no feedback to their contributions, they could disengage from the consultation. The satisfaction survey conducted by Cap Collectif[44] supports

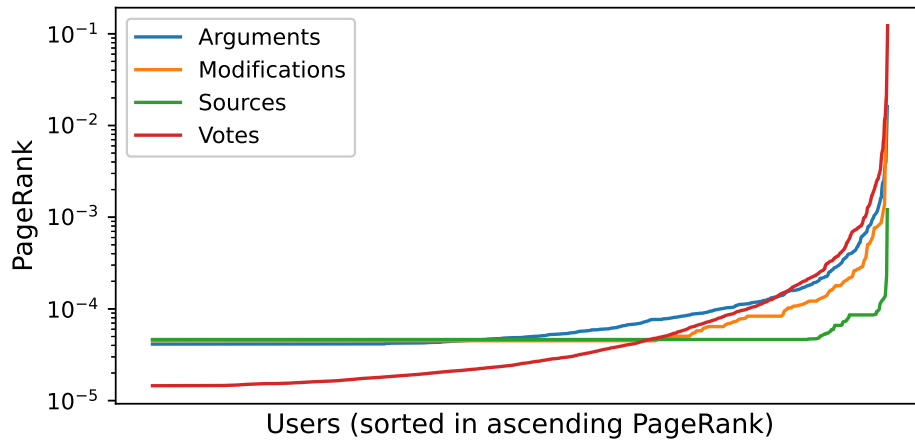


Figure 3.2: PageRank of users in ascending order.

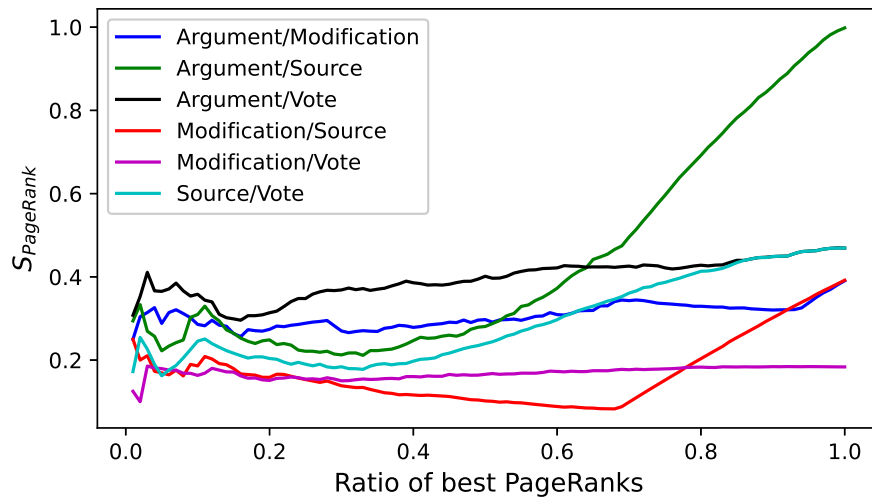


Figure 3.3: PageRanks correlation for different reply types.

this claim as it indicates that 81.4% of the respondents have visited the platform less than 5 times. This may also prompt users to change their behavior, from proposing new articles for a bill project to commenting and approving preexisting *Proposals*. This leads us to make the following recommendation:

REC. 1. *To give voice to a wider part of the audience, we recommend to highlight the newest Proposals, in a way similar to the chronologically sorted layout proposed by platforms such as StackOverflow, to improve the share of initial Proposals that receive attention and feedback. The RepNum platform already supports sorting the Proposals of a section chronologically; however, this is not the default selection.*

3.3.2 Analyzing the factors influencing the contributions

After obtaining the different PageRanks of the consultation, we investigate whether the highlighted centralization of contributions is directly related to contributors with

high PageRanks and their individual influence, or depends on the *Proposals* themselves, independently of their authors. To do so, we first investigate whether or not the PageRank of a user for a given type of contribution is correlated to their PageRank for the other types. This correlation allows us to know if the contributions of two different types are centralized around the same initial contributors, and therefore to identify if the users who tend to receive the most replies of a given type also receive the most replies of other types. Let:

- $type_a$ and $type_b$ denote types of contributions among $\{Argument, Modification, Source, Vote\}$;
- PR_{type_a} (resp. PR_{type_b}) denotes the list of users sorted by ascending order depending on their PageRanks for $type_a$ (resp. $type_b$) contributions –the list refers to all the unique users contributing to the consultation;
- $r \in [0, 1]$ denotes the share of users taken from the end of the lists PR_{type_a} and PR_{type_b} (i.e., $r = 0.05$ means we compute the similarity of the most active 5% contributors for $type_a$ and $type_b$ contributions); and
- $PR_{type_a,r}$ (resp. $PR_{type_b,r}$) denotes the sublists of PR_{type_a} (resp. PR_{type_b}) obtained after picking the r share of the highest PageRanks in both lists.

The similarity ratio $s_{PageRank}$ provides the level of similarity between $PR_{type_a,r}$ and $PR_{type_b,r}$ (i.e., at $s_{PageRank} = 0$, none of the users in $PR_{type_a,r}$ are in $PR_{type_b,r}$, and at $s_{PageRank} = 1$, both lists are identical):

$$s_{PageRank} = \frac{|PR_{type_a,r} \cap PR_{type_b,r}|}{|PR_{type_a,r} \cup PR_{type_b,r}|}$$

We compute the similarity ratio between two contribution types using the filtered PageRanks lists computed in § 3.3.1. The result depicted in Figure 3.3 shows that the degree of similarity remains under 0.5 for the low values of r ($r < 0.2$). For certain pairs of contributions (e.g., the *Argument/Modification* and *Argument/Vote* pairs), the similarity remains stable between 0.3 and 0.4. Consequently, we argue that PageRanks for these types of contributions are correlated. The correlation is weaker for other pairs (e.g., *Modification/Source* and *Modification/Vote*) that seem to be independent. The fact that the *Argument/Source* curve of the graph is the only one that converges towards 1 for the highest values of r is due to the fact that both *Arguments* and *Sources* can only be used as replies to *Proposals* and *Modifications*, i.e., $PR_{Argument}$ and PR_{Source} contain the same users after filtering. We can conclude that the users who receive most of the *Argument* replies tend to also centralize other types of replies, but this centralization does not apply for most other types.

Although contributors tend to weakly centralize contributions of different types in most cases, we want to know if this centralization is stronger around *Proposals*. This leads us to compute the similarity of the *Proposals* receiving the most replies of each type. Let:

- P_{type_a} (resp. P_{type_b}) denotes the list of *Proposals* in ascending order depending on the number of replies of $type_a$ (resp. $type_b$)

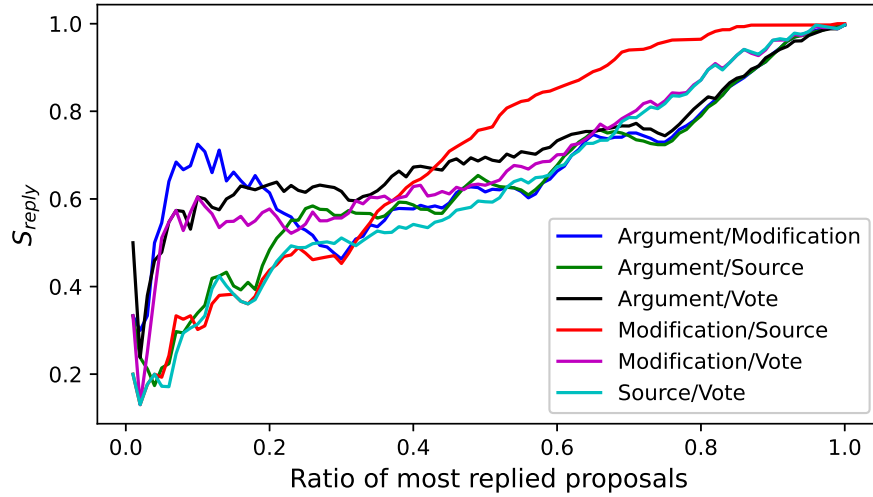


Figure 3.4: Correlation of the amount of replies to a proposal for different contribution types.

- $P_{type_a,r}$ (resp. $P_{type_b,r}$) denotes the sublist of P_{type_a} (resp. P_{type_b}) obtained after picking the r (r defined as before) share of the *Proposals* with the most replies.

The similarity ratio s_{reply} , which provides the level of similarity between $P_{type_a,r}$ and $P_{type_b,r}$, is defined as:

$$s_{reply} = \frac{|P_{type_a,r} \cap P_{type_b,r}|}{|P_{type_a,r} \cup P_{type_b,r}|}$$

The results of the computation (see Figure 3.4) highlight two main tendencies. First, for the lowest values of r (i.e., $r < 0.2$), the levels of similarity for pairs of reply types follow two tendencies. The *Argument/Modification*, *Argument/Vote* and *Modification/Vote* pairs have a high level of similarity ($s_{reply} > 0.5$) while this correlation is clearly weaker for all the pairs that contain the *Source* type of reply, i.e., the *Argument/Source*, *Modification/Source* and *Source/Vote* pairs. The computation becomes less relevant as r grows since P_{type_a} and P_{type_b} contain the same users, i.e., s_{reply} converges towards 1 for the highest values of r . In other words, although the *Proposals* that gather the most replies of a given type tend to also receive the most replies of other types, the correlation does not apply to the *Source* reply type that appears independent from the others. However, we must consider that this lack of apparent correlation may result from the low number of *Source* replies in the dataset (see Figure 3.2) and must not be interpreted as an inverse correlation or a decorrelation.

Second, the comparison of the filtered version of the correlation of PageRanks (see Figure 3.3) with the graph obtained from the computation of s_{reply} evidences that the similarity level grows faster for the latter. This difference in the shape of the curves may be related to a low impact of initial authors on the concentration of replies around a single contribution. Therefore, replies would be sent to a contribution due to its content rather than its author.

Our analysis shows that *Proposals* that receive the most replies of a certain type –excluding *sources*– tend to also centralize replies of the other types. This can lead to situations where a *Proposal* receives, e.g., a great number of *Votes* and very few




Figure 3.5: Correlation of activity for each contribution type.

Sources, becoming de facto a subject that both meets a great interest from the contributors and proposes very few material to form an opinion. We derive the following recommendation:

REC. 2. *We recommend to give the opportunity to users to add tags to Proposals such as “sources needed” or “lacking arguments” to attract potential contributors who may enrich their least complete parts. As part of our ongoing work, we are also investigating the automated labeling of contributions depending on the types of contributions they lack using a machine learning approach based on, inter alia, the number of replies of different types to a Proposal or their content, as proposed in [81].*

3.4 How to mitigate the influence of the most active users?

In order to further understand the behavior of contributors, we analyze the correlation between the activity of a user for different types of contribution. This correlation allows us to know whether the most active contributors for a type of contribution are also very active in another type, or the two behaviors are independent. Let:

- U_{type_a} (resp. U_{type_b}) be the list of users sorted by ascending order depending on the number of contributions of $type_a$ (resp. $type_b$) they have published;
- $U_{type_a,r}$ (resp. $U_{type_b,r}$) be the sublist of U_{type_a} (resp. U_{type_b}) obtained after picking the r (r defined as before) share of users with the highest number of contributions in the list.

The $s_{activity}$ ratio, which provides the level of similarity between $U_{type_a,r}$ and $U_{type_b,r}$, is defined as:

$$s_{activity} = \frac{|U_{type_a,r} \cap U_{type_b,r}|}{|U_{type_a,r} \cup U_{type_b,r}|}$$

Figure 3.5 highlights two different kinds of evolution of the similarity depending on r . The first group includes pairs of contribution types for which the level of similarity does not significantly increase while $r < 0.4$, and raises significantly afterwards. This group is composed of all the pairs containing the *Vote* contribution type. Here, the similarity level increases only because the wider both sublists are, the more likely a user is to be found in both of them. Therefore, we cannot declare that there is a correlation between the activity in *Vote* and other types of contributions.

The second group contains the pairs for which the level of similarity is low for $r < 0.1$ but drastically increases afterwards and finally remains almost equal to 1 after the first half of the graph, resulting in a square root-shaped curve. This group contains all the pairs containing two text-based contribution types, i.e., any type except *Vote*. For these pairs of contribution types, even if the 5-10% most active contributors are completely different, we can see that their top 20% is very similar. The most visible example of this group is the *Proposal/Source* curve of the graph. We conclude that, for pairs in this group, there is a correlation between the activity in both contribution types. This type of behavior is similar to the one identified by Vasilescu *et al.* [165] in Q&A platform StackOverflow and version control platform GitHub, with persons producing a high number of commits in GitHub projects also tending to answer a high number of questions on StackOverflow. However, we also identify that the most active contributors ($r < 0.05$) in one type tend to be specialized in this type and do not produce as many contributions of other types.

The analysis further indicates that, for text-based contributions, even though the 10 to 20% most active contributors for a specific type of contribution tend to also contribute to other types above the average, there is a small minority of specialized contributors who focus on producing contributions of one specific type. This leads us to infer that the presence of extremely active contributors, whether they are specialized in one type of contribution or versatile, can be a potential source of bias in participation. Indeed, these kinds of behavior can lead to situations where, for a given *Proposal*, most reply contributions of one or more types are submitted by a specific opinion group. This kind of situations could reduce the capability of contributors to form an informed opinion on a *Proposal* and bias their final (dis)approval. Consequently, we introduce the following recommendation:

REC. 3. *To reduce the risks induced by the centralization of contributions of a given type around a small set of contributors, it is necessary to identify the Proposals whose replies are made by a small set of active contributors. This detection can be achieved using graph-based anomaly detection [4] where: contributors and Proposals are two different types of vertices, and a contributor is linked to a Proposal by an edge if they have submitted a contribution related to the said Proposal. The Proposals that are detected as receiving an abnormally high number of replies from a few contributors are then highlighted to other contributors, so that they can provide new perspectives. This detection could lead to a reduction of the bias induced by the original contributors' perspective.*

3.5 Does the positioning of *Proposals* have an impact on participation?

In order to better understand which users and which *Proposals* tend to receive the most replies, we hypothesize that the positioning of *Proposals* on the consultation's user interface has an impact on their popularity. Indeed, the *RepNum* platform displays *Proposals* according to the two following criteria:

Categorization: *Proposals* are grouped according to the categories of the consultation's structure (see Figure 3.1). As a result, contributors have easier access to the content related to the categories at the top. This potentially creates a bias toward these categories and thus an impact on replies to content in lower categories.

Pinning: Within each category, the first *Proposals* displayed are always the ones composing the initial bill project, that is, the ones proposed by the government. These *Proposals* are “pinned” at the top of their category. The other *Proposals* can be sorted according to different criteria (e.g., date of publication, number of *Votes*, etc.) but are sorted randomly by default. This layout implies that the *Proposals* made by the government are more accessible to the contributors. This is a potential source of bias that can lead users to focus on government *Proposals* at the expense of public participation.

To assess the actual impact of this design choice of the *RepNum* platform, we statistically analyze the number of replies of each type to the *Proposals*, according to the two above criteria. In our first analysis, we group *Proposals* according to their categorization (see Figure 3.1). We then compute the quartiles and median for the number of replies of each type to a *Proposal* depending on the three levels (see Table 3.4 in Appendix). Regardless of the chosen categorization level, we notice small variations in the three statistical values. These variations tend to be wider for *Votes* than for other types of replies, with peaks identified in Section II.2.2, and Section III.3.1 of the consultation. However, we cannot identify any clear decrease or increase in the number of replies depending on the positioning of the category in the user interface. Nevertheless, we can see that the way categories are presented has a clear impact on the number of *Proposals* submitted within a section. Indeed, the number of *Proposals* per section tends to decrease as the category is positioned lower in the user interface. This tendency is more significant for coarser levels. Therefore, we consider that even though the impact of the positioning of contributions based on categorization is not significant enough to be identified as a bias generator in the contribution of replies to a *Proposal*, it has an impact on the number of *Proposals* submitted in a section and reduces the participation for the related subjects. We make the following recommendation:

REC. 4. *The number of Proposals related to the topmost categories – as displayed – being higher than that of the lower categories, we recommend either to sort categories randomly so that each contributor is exposed to different topics of the consultation, or to implement another method to display categories of Proposals on a more horizontal layout.*

The second analysis compares the popularity of all the *Proposals* of the consultation with those made by the government. First, we group the *Proposals* made by the

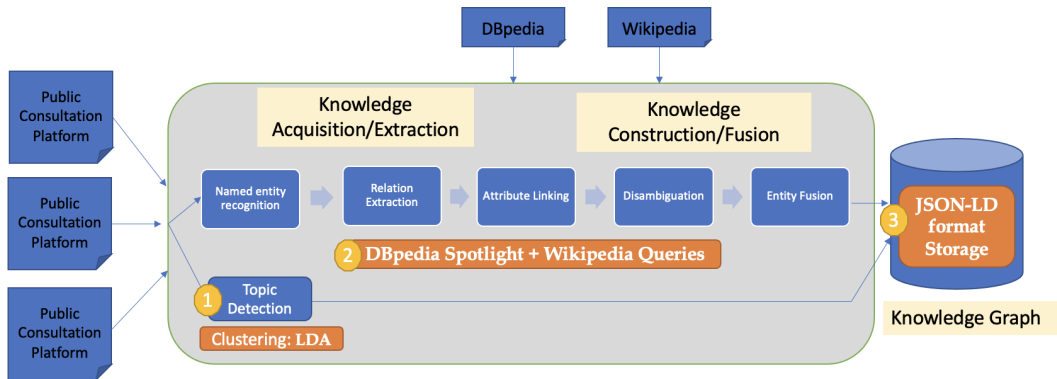


Figure 3.6: Steps for building the Online Public Consultation KG.

government according to their part, chapter and section (see Table 3.4 in Appendix). The number of *Proposals* from the government being particularly low (between 1 and 6 per section), we only compute the median for the number of replies of each type. Then, we compare these values to the associated medians and quartiles for the whole set of *Proposals* computed previously. The comparison shows that for every type of reply and category level, the median of replies to government *Proposals* of a category is always higher or equal to the third quartile of the number of replies to *Proposals* in this category. We conclude that pinning *Proposals* has a significant impact on the number of replies received, and consequently is a bias generator and centralizes the participation around the *Proposals* made by the government. This leads to the following recommendation:

REC. 5. *As the pinning of Proposals tends to centralize the replies around a small set of pinned Proposals, we recommend either not to implement this option or to change the selection criteria to prioritize pinning Proposals that would benefit the most from a high number of replies, such as those that are highly controversial.*

3.6 Building an Online Public Consultation Knowledge Graph

The three previous sections of this chapter allows us to identify biases in the design of the *RepNum* platform and to recommend solutions to avoid reproducing such issues in similar online consultations. In this section, we propose to improve information retrieval in a consultation platform using a knowledge graph. Figure 3.6 shows the different stages of our methodology for building the consultation KG. We detail these stages in the rest of this section by building a KG for the *République Numérique* consultation.

3.6.1 Topic modeling

We use Latent Dirichlet Allocation (LDA) [79] on the proposals submitted to the *RepNum* consultation to identify the different topics. Each proposal is represented as a document, and LDA is an algorithm for topic modeling in a set of documents

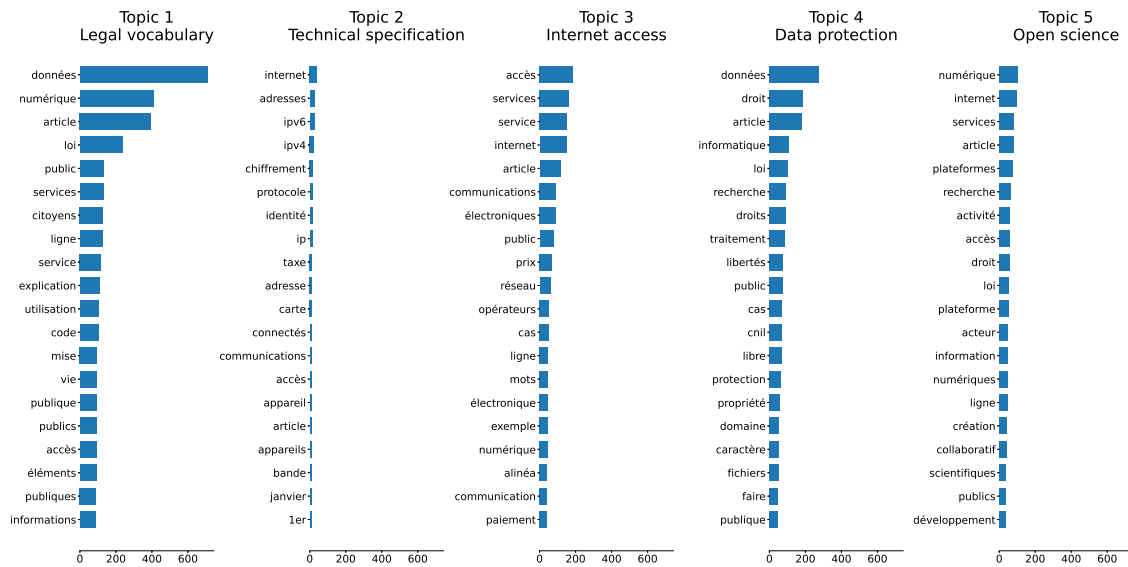


Figure 3.7: LDA representation of the main topics discussed in the *République Numérique* citizen consultation and their most frequent words. Topic labels were chosen a posteriori by the authors of this work to help differentiate them.

where each topic is associated with the different N-grams² based on their number of occurrences in the different documents. Each proposal is then associated with all relevant topics based on a truth value score. Using this algorithm, we need to choose the number k of topics to identify. We applied LDA for $k \in [4, 15]$ and assessed the suitability of each version of the topic modeling based on the consistency of the 20 most frequent words for each topic. We consider $k = 5$ to be the parameter giving the most coherent results for the *République Numérique* citizen consultation. Figure 3.7 shows the five main topics identified and their most frequent words. We use these topics, in addition to the proposal’s title, author, and content, to further characterize the different proposals.

3.6.2 Entity Recognition and Linking

We use a French-language adaptation of DBpedia Spotlight [49, 108] (DBS for short in the following) to identify entities in our proposals. DBS extracts important words from a given text and links them to resources in DBpedia [18]. We have set the “confidence” parameter to 0.6 and left everything else to default. Our observations about DBS concern the French model. Linking the most important elements in the proposals to DBpedia resources allows us to associate these elements with definitions from Wikipedia and connect proposals based on the resources they mention. To eliminate certain common false associations identified in preliminary tests, we filter the obtained results:

- DBS tends to identify years and dates as important elements of the proposals and link them with the aggregation Wikipedia page for events happening at this given date - e.g., https://en.wikipedia.org/wiki/December_13-. In

²A sequence of 1 or more words usually written consecutively. For example, both “text” and “knowledge graph” are N-grams in the context of this section.

the context of an online consultation, such elements are often mentioned in contexts where this linking is not relevant. Therefore, we exclude these resources before performing Named Entity Linking (NEL).

- We also observe that the model tends to identify certain N-grams as film titles. This has almost only led to incorrect identification of common N-grams as film titles in the context of the *RepNum* consultation. Consequently, we have filtered non-documentary films out of the resource graph with which we associate our entities. The choice to keep documentary films is made to give citizens the possibility to add references to support their proposals.

Then, we identified an additional issue: acronyms with multiple definitions are often associated with the wrong resource. For example, the acronym “CADA” is used in a proposal to mention the “Administrative Documents Access Commission” (in French, *Commission d’Accès aux Documents Administratifs*). However, this entity is linked to “Shelter for Asylum Seekers” (in French, *Centre d’Accueil de Demandeurs d’Asile*), even though “administrative documents” is mentioned further in the proposal. This issue has been confirmed with similar results by using DBS on another online consultation called RUA (Open data available at <https://www.data.gouv.fr/fr/datasets/consultation-vers-un-revenu-universel-dactivite-1/>). We solve this issue by applying the following process to each acronym identified as important in a proposal.

- If DBpedia has a node with a name identical to the acronym:
 - If this node is a specific resource (e.g., a Wikipedia page) or redirects to it, we associate the acronym with this resource.
 - Else, if the node is a disambiguation node for multiple resources, we gather the textual abstract attached to each resource. Then, using the SBERT [137, 139] model for Semantic Textual Similarity mining, we apply cosine similarity to each abstract with the proposal containing the acronym. Finally, we associate the acronym with the resource whose abstract has the highest cosine similarity with the proposal.
- Else, using the Wikipedia API, we search for the acronym and gather the 5 best results. Then, similarly to the previous dash, we use SBERT’s cosine similarity to identify the Wikipedia page whose abstract is the most similar to the proposal. Finally, we link the acronym with the DBpedia resource associated with the identified Wikipedia page.

Performing Named Entity Recognition (NER) and NEL as mentioned earlier generates a set of proposals where certain N-grams are associated with a DBpedia resource. The obtained results are presented in Table 3.3. A considerable proportion of the proposals (16.2 %) are not associated with any entity, likely because these proposals are shorter than others. Specifically, while the median number of tokens³ for proposals with zero annotations is 80, it is 188 for proposals with one or more annotations. Moreover, more than one-third of the proposals (36.6 %) are associated

³In NLP vocabulary, a token refers to a part of a string. In this study, a token is either a number, a word, or a punctuation mark.

Table 3.3: NER and NEL results using DBpedia Spotlight on proposals of the RepNum consultation.

Metric	Value
Number of proposals	692
Raw number of entities	3,463
Number of different entities	1,015
Number of entities linked with exactly 1 proposal	680
Number of entities linked with 5 or more proposals	82
Number of annotations for entities linked with 5 or more proposals	1,655
Number of proposals without any annotation	112
Number of proposals with 5 or more annotations	253
Median number of annotations per proposal	3

with at least five entities. Also, the number of entities linked to at least one proposal is greater than the number of proposals. However, although the number of entities connected to five or more proposals is significantly lower, almost half (47.8 %) of the annotations link proposals to these entities.

3.6.3 Resulting Knowledge Graph

Figure 3.8 shows a visual representation of our KG. We store the generated KG using the JSON-LD format [157]. As JSON-LD is an RDF serialization, querying the graph using SPARQL is possible.

The graph we propose facilitates the issuance of new queries in consultations by both citizens and platform administrators. For instance, one can retrieve all proposals mentioning a specific entity while focusing on a particular topic, such as “all the proposals mentioning the National Commission on Informatics and Liberty (CNIL)” with an LDA-score greater than 0.6 for the “Data Protection” topic. Furthermore, the model allows for summarizing contributions by identifying the topics and entities mentioned in a consultation or linking entities to specific groups of users. By using this method, platform administrators can enhance the transparency of their reports on citizen consultations by incorporating graph-based data alongside responses to a diverse array of proposals⁴, selected either randomly or based on predetermined criteria. Lastly, our work can be generalized to most online participatory platforms to connect contributions from various online consultations.

3.7 Threats to validity

For the analysis of the centralization of contributions around contributors and *Proposals* (§ 3.3.2), we did not take into account the relevance of contributions. Indeed, certain contributions can be out of the scope of the consultation, mischievous, or duplicates of pre-existing *Proposals*. This factor could be meaningful to explain why a *Proposal* receives many replies or not, as the least relevant *Proposals* are likely to

⁴For instance, see <https://www.republique-numerique.fr/project/projet-de-loi-numerique/step/reponses>

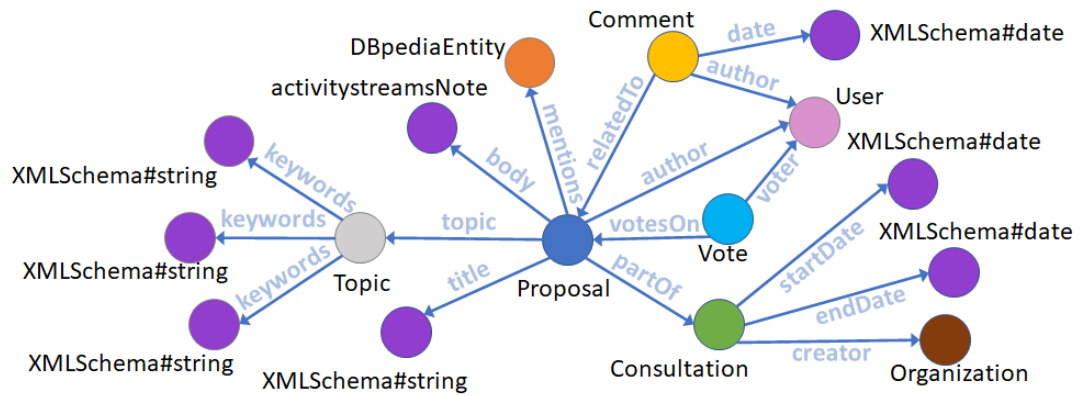


Figure 3.8: Visualization of a generic consultation stored in JSON-LD.

be ignored by contributors. This should be considered to balance the contributions received by *Proposals*, so as to not highlight irrelevant *Proposals* that receive little attention.

We focus this study on a specific consultation. It is therefore subject to a demographic bias concerning the studied population and our findings could differ in other consultations. We mitigate this risk by choosing the largest dataset we have found, concerning the number of both contributors and contributions. However, the dataset still contains very few entries for certain types of contributions such as *Sources*, which reduces the representativeness of our sample for these types of contributions. Therefore, reproducing similar work on different datasets is worth investigating to generalize our recommendation to other participatory platforms.

When we evaluated the impact of the positioning of *Proposals* on participation (§ 3.5), our conclusions and recommendations were based on the default positioning of *Proposals* on the user interface. However, we were not able to get statistics about the percentage of contributors using this default interface. Still, lacking this information has no impact on the validity of REC. 4 since, regardless of the modifications of the interface proposed by the platform, categories are always displayed in the same order. Nevertheless, if most users did choose a non-default positioning of *Proposals* which would not put pinned *Proposals* forward, then the centralization of contributions around the said pinned *Proposals* would be only caused by the fact that they have been issued by the government. In that case, REC. 5 would not effectively solve the identified problem since it would not be induced directly by the user interface.

3.8 Related Work

Ben Jabeur *et al.* [24] introduce three new algorithms based on PageRank to characterize the specific behavior of three different groups (*influencers*, *leaders* and *discussers*) of key users in Twitter. Weng *et al.* [174] present another PageRank-based algorithm, also targeting Twitter but focused on the detection of influencers depending on the topic. Our work distinguishes from the above studies by addressing participatory platforms. However, their analysis identifies different types of following behaviors depending on the topical similarity between users, in a way similar to ours. This work suggests that the dedicated customization of PageRank would

allow the detection of influential contributors according to the given use of online citizen participation platforms, which is area for our future work.

Conover *et al.* [46] propose to use label propagation [136] to identify clusters of users belonging to the same opinion group. Their approach allows to identify the patterns of interaction between users from the two clusters they identify –left- and right-leaning users– for the different types of interactions proposed by the platform. This method can be relevant in our case, especially as a mean to improve the reliability of REC. 3 in the detection of opinion groups centralizing the discussion around a specific proposition.

Khan *et al.* [85] present a requirement engineering approach based on the extraction and evaluation of arguments from end-users feedbacks. They apply an argumentation theory and machine learning-based approach to establish requirements for a Google Maps feature. Our approach differs greatly from theirs since we draw our recommendations based on users’ behaviors and on the flaws they highlight on the platform we study. We argue that our approach has the advantage of directly relying on users’ interactions with the platform rather than on users feedback. Therefore, our approach allows the identification of issues even if users are not aware of their existence. However, argument mining-based approach has the advantage of directly giving voice to platform users and therefore could be used to improve online participatory platforms based on arguments provided by citizens.

While the focus of this chapter has been on informing the design and/or enhancement of participatory platforms based on the analyses of a past consultation, related studies are concerned with questioning the overall development process of these platforms. For instance, Knutas *et al.* [86] present a case study of the software development processes that emerged from the “*Code for Ireland*” initiative. They provide a description of the implemented practices together with their achievements. They also identify specific challenges in the development process, notably concerning the specification of requirements or the engagement of stakeholders, and propose a set of recommendations to solve the problems identified.

Finally, the study of the *RepNum* consultation has received much attention from the French-speaking social science community. This includes the work of Laurent *et al.* [91] who analyze the graphic design of the *RepNum* user interface and the interactions it enables with the platform. They highlight the inadequacy of the contribution means offered by the platform with respect to the diversity of the points of view expressed. They also investigate the use of the platform in the context of the French representative democracy. Alexis *et al.* [5] also analyze the functionality of the *RepNum* platform, taking into account the tools offered to the user for the consultation and investigating the types of discourse induced. They highlight a depoliticization of the discourse through the neutralization⁵ of the subject. These studies are complementary to ours and discuss other keys to understand how the *RepNum* consultation worked.

⁵Erasure of oppositions between different points of view and the production of signs of the evidence of a single point of view, assumed transparent (translated from [122]).

3.9 Conclusion

We analyzed the different flaws of the *RepNum* platform regarding citizen participation. Our findings show a high degree of centralization of contributions around a small set of *Proposals*. This centralization is more related to the content of *Proposals* themselves than to the identity of their authors. We also highlight the existence of extremely active contributors, with some of them being specialized in a specific type of contribution and others being more versatile. Finally, we identify the impact of the positioning of *Proposals* on the number of related contributions submitted. These different situations constitute threats to participation, and therefore to the inclusion of citizens in public decision making. To solve the problems arising from them, we formulate five recommendations oriented towards the design of participatory platforms.

We conclude that the design of participatory platforms should be considered as a vector of bias to the participation itself, and the behavior of contributors is a meaningful indicator to identify such issues. However, to the best of our knowledge, contemporary participatory platforms do not implement certain modern software features, which could be used to improve their existing features, involve citizens deeply in the process and respond to new use cases. We propose implementing a tool to improve exploration of a set of proposals using an automatically constructed KG, to facilitate both the identification of relevant proposals for a citizen and the summarization of contribution by platforms administrators. Moreover, modern CSCW tools such as collaborative editing would help participants elaborate collective proposals. Social recommendation [60] could also be leveraged to highlight relevant topics for contributors or help them find co-writers for common proposals.

The next chapter of this thesis will present the second contribution of my Ph.D.: an interview-based study presenting a description framework for online citizen participation based on the objectives pursued by platform maintainers and the technical features implemented.

Appendix: Data tables for Section 3.5 (Also available on the Git repository)

Table 3.4: Amount of replies to contributions of all types for each category of the consultation.

Part	Chapter	Section	Nb. of Proposals	Arguments			Modifications			Sources			Votes		
				Q1	Med	Q3	Q1	Med	Q3	Q1	Med	Q3	Q1	Med	Q3
Part I	Chapter 1	Section 1	160	1	2	5	0	0	1	0	0	0	14	24	81
		Section 2	47	1	2	5	0	0	1	0	0	0	7	23	48
		Section 3	33	1	3	6	0	0	1	0	0	0	9	23	75
		Whole chapter	240	1	2	5	0	0	1	0	0	0	11	23	76
	Chapter 2	Section 1	59	0	2	5	0	0	0	0	0	0	10	24	52
		Section 2	25	1	2	6	0	0	1	0	0	1	8	26	203
		Whole chapter	84	0	2	5	0	0	1	0	0	1	10	24	81
Whole part		324	1	2	5	0	0	1	0	0	0	10	24	77	
Part II	Chapter 1	Section 1	30	1	4	6	0	0	1	0	0	0	11	24	62
		Section 2	19	0	1	2	0	0	1	0	0	0	4	14	21
		Section 3	23	1	2	6	0	0	0	0	0	0	8	15	29
		Section 4	31	0	1	4	0	0	0	0	0	0	8	12	39
	Whole chapter	103	1	2	6	0	0	1	0	0	0	8	16	48	
	Chapter 2	Section 1	82	0	1	3	0	0	0	0	0	0	4	9	27
		Section 2	19	0	2	10	0	1	1	0	0	1	8	42	248
Whole chapter		101	0	1	4	0	0	1	0	0	0	4	11	43	
Whole part		204	0	2	5	0	0	1	0	0	0	6	13	48	
Part III	Chapter 1	Section 1	60	0	1	3	0	0	1	0	0	0	3	8	30
		Section 2	30	1	2	5	0	0	1	0	0	0	4	10	28
		Whole chapter	90	0	1	4	0	0	1	0	0	0	4	9	28
	Chapter 2	Section 1	19	1	3	5	0	0	1	0	0	1	5	13	37
		Section 2	12	0	3	5	0	0	0	0	0	0	4	25	83
		Whole chapter	31	1	3	5	0	0	1	0	0	1	5	16	83
	Chapter 3	Section 1	13	2	3	6	0	1	1	0	0	0	23	114	396
		Section 2	14	0	3	4	0	0	1	0	0	0	3	13	76
		Section 3	16	1	4	10	0	0	3	0	0	0	12	52	174
		Whole chapter	43	1	4	6	0	0	1	0	0	0	12	51	185
Whole part		164	1	2	5	0	0	1	0	0	0	4	13	70	

Table 3.5: Amount of replies to governmental contributions of all types for each category of the consultation.

Part	Chapter	Section	Nb. Proposals	Med. Arguments	Med. Modifications	Med. Sources	Med. Votes
Part I	Chapter 1	Section 1	3	44	32	3	1319
		Section 2	1	38	30	3	892
		Section 3	3	23	19	1	777
		Whole chapter	7	38	27	1	892
	Chapter 2	Section 1	1	25	49	2	822
		Section 2	2	40	12	1	706
		Whole chapter	3	40	49	2	822
Whole part		10	38	27	1	822	
Part II	Chapter 1	Section 1	1	40	24	1	1426
		Section 2	1	47	38	0	796
		Section 3	2	24	16	1	379
		Section 4	1	31	24	2	488
		Whole chapter	5	38	24	1	669
	Chapter 2	Section 1	6	23	16	0	495
		Section 2	1	35	26	1	853
Whole chapter	7	25	26	1	561		
Whole part		12	28	24	1	561	
Part III	Chapter 1	Section 1	1	22	41	2	368
		Section 2	2	9	7	5	226
		Whole chapter	3	13	7	5	368
	Chapter 2	Section 1	1	21	10	5	392
		Section 2	1	81	30	6	1380
		Whole chapter	2	21	10	5	392
	Chapter 3	Section 1	1	19	24	7	498
		Section 2	1	45	22	5	408
		Section 3	1	24	11	5	833
		Whole chapter	3	24	22	5	498
Whole part		8	21	11	5	392	

Chapter 4

Objectives and features of contemporary platforms

Contents

4.1	Introduction	69
4.2	Methodology	71
4.3	Different objectives of citizen participation platforms	73
4.3.1	Surveying citizens' opinions	73
4.3.2	Dialogue between citizens and an institution	74
4.3.3	Creation of opinion groups or neighborhood councils	75
4.3.4	Co-construction of projects	76
4.4	Different types of platforms for contemporary online citizen participation	77
4.4.1	"Suggestion box"-like platforms	77
4.4.2	Interactive platforms	79
4.5	How to enrich online participatory platforms	80
4.5.1	Collaborative systems	80
4.5.2	Natural Language Processing for large-scale processing of contributions	81
4.6	Threats to validity	83
4.7	Conclusion	83

This chapter is based on our paper "Citizen participation platforms' objectives and how to reach them: an interview-based analysis", written with Rafael Angarita and Valérie Issarny. At the time of submission of my Ph.D. thesis, this paper has not been published yet.

4.1 Introduction

The use of online platforms supporting citizen participation has become more common in the recent years, allowing the organization of citizen participation at larger

scales [12, 91], both concerning their number of contributors and the geographical distance between them. Such platforms implement a wide range of features to help citizens contribute to a participatory project. These features implement the rules of a given participatory process as defined by the organizing institution. They also define the degree of possible interaction between the contributors of a given platform, allowing, for example, the creation of altered versions of a pre-existing contribution. However, Ianniello *et al.* [75] identify that bad quality participatory process design or citizen collaboration has been a major obstacle to successful citizen participation processes in previous contexts. Chapter 3 also identified flaws in the design of a specific online consultation and its human-computer interface, leading to measurable biases in the contributions produced by citizens. Therefore, we consider that, in the context of online participatory platforms, the features implemented are a key factor to the success of citizen participation initiatives. Moreover, the high number of contributions have also been identified as the origin of an information overload experienced by citizens [38, 132]. Arana-Catania *et al.* [13] propose implementing machine learning-based features to reduce this information overload and help contributors identify relevant contributions.

Apart from the features they implement, participatory platforms also differ depending on their objectives. One way to characterize the objectives of a participatory process can be to use Arnstein’s ladder of citizen participation [15]. The ladder introduces eight levels, with each corresponding to a different level of citizen empowerment, from manipulation to citizen control (see Figure 4.1). Each level implies different forms of interaction between citizens and the concerned organization as well as among the citizens themselves. However, even though the level of citizen empowerment is crucial to differentiate participatory platforms, the objectives of citizen participation initiatives can also be defined depending on the type of output expected by the institution. This differentiation overlaps with the one proposed by Arnstein’s ladder since the higher degrees of the ladder – called “degrees of citizen power” – cannot be achieved if the expected input is simple and individual (e.g., answering a survey). Moreover, while platform administrators are not always aware of the specificities of each degree of participation, they clearly identify the type of output expected in the platforms they maintain and therefore can be more reliably trusted to describe outputs produced by contributors than the positioning of their platform on Arnstein’s ladder. Consequently, we consider that the type of output expected is a more suitable criterion to differentiate the objectives of participatory platforms than their degree of citizen participation.

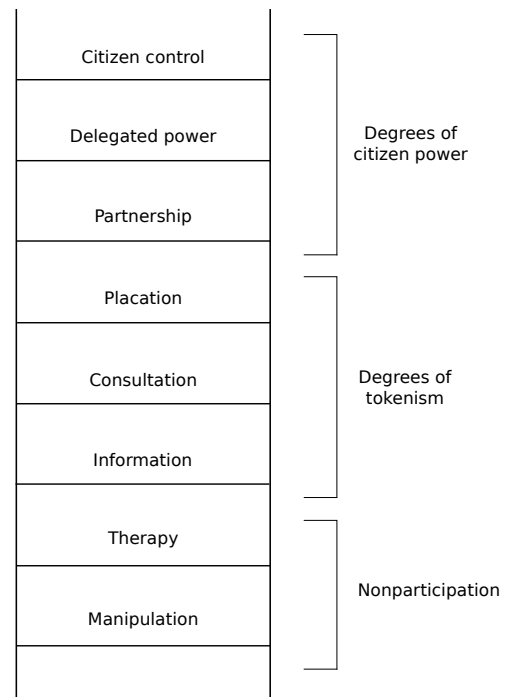


Figure 4.1: Arnstein’s ladder of citizen participation [15]

Using the two aforementioned axes to differentiate participatory platforms –

objectives pursued and features implemented –, this chapter answers the following research questions:

RQ1: What are the different objectives set by citizen participation platforms maintainers¹?

RQ2: How do the features implemented in participatory platforms impact the design of citizen participation initiatives?

To answer these questions, we rely on a series of five semi-directed interviews with citizen participation platforms administrators and designers. These interviews allow us to differentiate platforms depending on their objectives (§ 4.3) and on the means they implement to accomplish these objectives (§ 4.4). To ensure the representativeness of our panel of interviewees, we chose to interview stakeholders with diverse professional profiles, from cities of different sizes and who use different software for their participatory platforms (§ 4.2). Based on the different objectives and features identified from our interviews, we formulate recommendations to improve the experience of contributors and administrators of citizen participation platforms (§ 4.5), as well as to guide platform administrators and designers on the development of future participatory platforms. These recommendations are focused on machine learning-based identification of relevant contributions, collaborative editing of contributions, and the integration of pre-existing features that have yet to be implemented within the context of citizen participation platforms.

4.2 Methodology

To better understand users' behaviors in online citizen participation platforms and the specific rules of the aforementioned platforms, we conducted a series of five semi-directed interviews with citizen participation platform administrators and providers (profiles of the interviewees are listed in Table 4.2). The choice of using interviews rather than empirical data produced by contributors of citizen participation platforms is motivated by our need to identify the objectives pursued by the organizers of participatory processes (e.g., cities, and public institutions). We consider that such information cannot be obtained without directly interviewing stakeholders working for the aforementioned organizing bodies. The questions we asked were related to the objectives they seek to achieve, the functionalities implemented and the means used to gather and synthesize contributions (see Table 4.1). These interviews allowed us to identify different use cases for citizen participation as well as the software tools implemented in contemporary platforms.

We consider that our panel of interviewees quite accurately represents the diversity of citizen participation platforms. Indeed, our interviewees work for city of various sizes with two under 100k inhabitants, one over 200k and one over three million, in different geographic locations and at different degrees of technical implication. The platforms our interviewees work for have different purposes, including

¹In this chapter, we differentiate platforms administrators and maintainers. We use “administrators” to reference the persons who manage participatory platforms and summarize contributions produced. The word “maintainer” is used to reference the entity which creates and maintains a participatory initiative taking place in an online platform (e.g., an institution).

Table 4.1: Questions for the interviews. The questions were modified in the case of Interviewee #4 as their company did create and manage multiple participatory platforms.

#	QUESTION
#1	What is the main objective which motivated the creation of your participatory platform?
#2	What functionalities have been implemented in your participatory platform?
#3	What objective should the different functionalities implemented fulfill?
#4	What service of your institution is assigned to the management of your platform? Do you collaborate with an external company?
#4'	If you work with an external company, what kind of specifications did you provide? Technical or functional?
#5	How do you manage the platform and gather the contributions received?

Table 4.2: Profiles of interviewees. Population of their respective cities were reported in September 2022.

INTERVIEWEE #	PROFILE	CITY	POPULATION	PLATFORM
#1	Platform administrator and designer	Strasbourg	287 k	Created in-house
#2	Platform administrator	Antibes	73 k	Decidim
#3	Platform administrator	Issy-les-Moulineaux	68 k	Cap Collectif
#4	Platform-making company executive	N/A	N/A	Cap Collectif
#5	Technical Advisor	Madrid	3.3 M	CONSUL

the collective elaboration of complex projects for the development of a city, the publication of suggestions to improve the quality of life in a neighborhood, or the creation of citizen initiatives to work collectively on a given topic. The means of contribution offered to citizens by the platforms managed by our interviewees represent a large array of possibilities, from the sending of ideas to the organizing institution to the collective elaboration of projects. Moreover, the platforms of our panel are developed both using software explicitly created for the occasion and three different all-purpose participatory systems – namely Decidim, Cap Collectif and CONSUL.

Finally, our interviewees have different types of involvement in their respective participatory platforms. Interviewee #1 is the lead designer of the platform developed for their city, their role is to propose and develop new features for the platform based on the specifications proposed with citizens. Interviewees #2 and #3 are administrators of their respective platforms, they monitor the activity on the platforms, summarize the contributions produced by citizens and interact with neighborhood councils to develop the ideas proposed. Interviewee #4 is an executive for a company designing participatory platforms and therefore has a broader view on the different objectives of their customers. Their company has created participatory platforms for cities of various sizes, colleges, companies or national consultations. Interviewee #5 is a technical advisor for citizen participation in their municipality, and therefore is involved in the technical evolution of their city’s participatory platform.

All interviewees were located in France or Spain and were remotely interviewed by video conferencing. Video recordings were collected with interviewees’ consent and then transcribed by Rafael Angarita and myself. Quotes from the interviews in the following sections have been translated by Rafael Angarita and myself.

Moreover, we have also studied the functioning of different consultations that used these software platforms and where text-based contributions were possible (e.g., <https://www.consultation-rua.gouv.fr/>, <https://www.decidim.barcelona/>),

Table 4.3: Categorization of platforms maintained by interviewees according to their objectives

INTERVIEWEE #	PLATFORM OBJECTIVES
#1	Multiple use-cases in the same platform. All the objectives identified are implemented in at least one use case.
#2	Dialogue between the citizen and an institution and creation of geographical communities based on the neighborhood councils system.
#3	Surveying citizens' opinions and dialogue between the citizen and an institution. In certain cases, proposals are then transferred to neighborhood councils.
#4	Multiple platforms pursuing all the objectives identified.
#5	Surveying citizens' opinions, dialogue between the citizen and the municipality and co-construction of projects.

as well as the resulting contributions obtained. We chose these consultations for being those with the highest number of contributions for their respective software platforms while still having their platform or open data available online at the time when we made our selection (September 2021). Certain examples cited in the rest of the chapter have been taken from these consultations but they are considered as background knowledge for this chapter. While none of these platforms contradict the conclusions of this chapter, we primarily base our work on the series of interviews mentioned earlier.

4.3 Different objectives of citizen participation platforms

One of the key questions in our interviews was about the objectives of citizen participation platforms (Question 1, Table 4.1). We identify four categories of objectives for platform maintainers: survey opinions, dialogue with the citizens, replicate offline participation groups and co-construct projects. These objectives are non-exclusive and all the platforms of our panel fit into at least two of these categories. We have compiled a list of the platforms that our interviewees work on and categorized them in Table 4.3.

4.3.1 Surveying citizens' opinions

Certain platforms of our panel are built to survey the opinions of citizens on several topics. Here, we limit this objective to the case where contributions of multiple contributors are gathered and summarized by the institution as a source of data concerning the general opinion of citizens, without the ambition to answer to contributions individually. Generally, these platforms implement a workflow where citizens answer a series of multiple-choice questions or short open-ended questions (e.g., "How would you beautify your neighborhood?"²). Interviewee #4 describes these platforms as follows:

[They are] opinion polls or surveys where you want the opinions of people, where you potentially want to put them in relation to their situation, where you can agglomerate data and analyze them. However, you do not

²Question found in the <https://idees.paris.fr> platform.

have any ambition concerning user interactions, like making everyone participate in the same space, take one another's opinions into account and interact together.

Here, the platform maintainers delegate to citizens either the ideation process, as described by [128], or the assessment of ideas already expressed in another context (e.g., offline neighborhood council, department of the city hall). Therefore, administrators of this type of platform generally argue that collaboration among citizens is unnecessary in these specific situations. Interviewee #1, who is the administrator of a platform where multiple participatory goals coexist, explains that they “try to differentiate the information depending on the degree of implication of the citizens and of feedback to them”. Use cases where the objective is to survey citizens' opinions are therefore annotated as the ones where contributors will receive very little individual feedback and will have very limited involvement in related projects. We identify this objective as closely tied to Arnstein's [15] definition of the “consultation” degree of participation, “inviting citizens' opinions”, with the already identified issue that “it offers no assurance that citizen concerns and ideas will be taken into account”.

4.3.2 Dialogue between citizens and an institution

An objective often mentioned by participatory platforms administrators we have interviewed is the capability to create a dialogue between citizens individually and an institution. Indeed, to a certain extent, all the platforms in our panel seek to facilitate such dialogue. Interviewee #2 even cites this objective early in our interview as their main motivation to create an online participatory platform in addition to their already existing offline neighborhood councils, since the latter “do not give the opportunity to everyone to express themselves and to stay aware of everything happening in the city”. We break this objective down into two sub-objectives: the capability for the institution to share information with the citizens – which is close to the “information” degree of Arnstein's ladder [15] – and the possibility for the citizens to provide feedback to the institution on specific topics – similarly to the “consultation” degree –, generally implemented using a two-way vertical workflow with little to no interaction among citizens. Moreover, our interviewees often quote this objective when talking about the signaling of simple issues within the city (e.g., garbage on the sidewalk, out-of-order traffic lights, etc.) or in the case of petition hosting systems, where the range of interactions between citizens tends to be very limited. Nevertheless, given a scope where a single citizen is not enough to provide complete feedback, collaboration among contributors can be useful to ponderate the opinions of citizens depending on their situation. For example, Interviewee #2 mentions the case of a problem with the road system which happened before they deployed their participatory platform:

Some people were telling us about a speed problem on a street and were asking us to make it a one-way lane. The mayor observed that people in neighborhood councils may not have the same concerns as the residents of the street, even though they live in the same neighborhood. [...] We gathered the opinions of the residents of this street by putting a letter

in their mailbox, and we realized that they did not fully agree with the proposals made by the neighborhood council.

In that case, the lack of a mechanism for citizens to review and engage with one another's proposals could have led to a situation where a decision is taken without the opinion of the most impacted residents. Even though this potential issue was avoided by the active intervention of the institution, we argue that allowing citizens to express their point of view in an online participatory platform would have also been a reliable mean to receive more diverse feedbacks in addition to those expressed by the members of the neighborhood councils.

This objective of dialogue is often pursued on a wider scope than only participatory platforms, generally including interaction between the citizens and institutions on social media. In the case of the platform maintained by Interviewee #3, their objective is to integrate the feedback received through social media to their "Citizen Relationship Management" system. This choice of multiplying the interaction channels, rather than centralizing all the means of information and feedback in the participatory platform, is motivated by their observation that "people tend to use any medium to communicate with the city". We notice that this diversity of channels is only mentioned by our interviewees when they talk about the objective of dialogue between citizens and institutions. We therefore consider that this need of multiple channels of participation is only addressed for this objective, even though it is relevant for the other objectives identified.

4.3.3 Creation of opinion groups or neighborhood councils

The possibility for citizens to constitute groups depending on different characteristics (e.g., interests, place of residence, etc.) is often cited by our interviewees as one of their objectives. These groups are typically communities of people producing common proposals – at the demand of an institution or at their own initiative – on a given topic. In the case of Interviewee #2, the objective of their platform was to "mimic the functioning of [their] neighborhood councils", i.e. to give the opportunity to citizens to enter a local group to submit proposals on the development of a given area. By placing emphasis on existing types of groups, such as neighborhood councils, the potential for contributors to form associations based on alternative criteria (e.g., interest or expertise in a given topic) is diminished. Interviewee #5 explains that "in Madrid, there was already citizen participation but it was an offline one and more oriented to already defined groups and associations". The mean they chose to allow the expression of a more diverse panel of contributors was "to let citizens participate individually", but other possibilities have been explored, in particular by Interviewee #1's platform:

We have a system to host citizen initiatives on our platform. Let's say I am a citizen, I have an initiative and I want to gather other citizens around my initiative. I can create an "embassy" on the city's participatory platform to find people. The city advertises the initiatives, we present them during a citizen participation council and people can decide that an initiative will be supported by the city.

The expected output of these platforms is not by itself a contribution from citizens to policies but the creation of groups of citizens according to their topics of

interest or place of residence. Consequently, we cannot identify a clear equivalence between the creation of these group and a specific degree of Arnstein's ladder. However, the acknowledgment of the existence of these groups allows citizens to produce common contributions to online consultations in a collaborative workflow. This specificity is important since the different platforms of our panel implement scarcely any means for online collaboration between contributors. The creation of groups inside participatory platforms can therefore be seen as a first step towards online collaboration in citizen participation, even though this collaboration still takes place outside of the platforms (e.g., in online CSCW platforms or in-person meetings). Moreover, the acknowledgment of the existence of groups working on a given topic allows institutions to include these groups in potential policy-making related with their topics of interest. For example, in Interviewee #1's platform, even though the city had created a citizen participation council to "review the tools for interacting with citizens [...] and validate the guidelines of the platform", an independent group was created by citizens to "define citizen participation" at the scale of their city. This group organized local events and eventually did produce a list of recommendations addressed to the municipality.

4.3.4 Co-construction of projects

The last objective we identify in participatory platforms is the possibility for citizens and institutions to elaborate common projects. The implication of both sides can greatly vary among platforms implementing this objective, from the possibility given to citizens to propose modifications to a project already conceived by an institution, to the collective elaboration of a project by representatives of the institution and citizens. Theoretically, this objective would therefore correspond to the degrees of participation between "consultation" and "partnership" inclusive [15]. However, in practice, we could not identify in the platforms of our panel an example of online co-construction where citizens were given actual decision power on the final project, which would consequently categorize the implementations of this objective in the "consultation" level of the ladder. This vertical organization, as well as the objective of allowing citizens to have an input on institutional projects, is described by Interviewee #4 as follows:

The common objective [of our platforms] is to allow the participation of citizen in public decisions making. [...] They do not only aim at being a space for expression but, most importantly, a place where we give visibility to a full cycle of a project in which participation is involved and also where we can organize feedbacks to participants by showing precisely which of their proposals were included or not to the final project.

This objective is the one where interaction between contributors is the most frequent. Indeed, platforms which aim at allowing co-construction of projects generally implement a horizontal workflow where citizens – and in certain cases, institutions representatives – are allowed not only to submit proposals, but also to interact with one another's proposals through comments or revised versions. However, the platforms of our panel hardly implement any means for online collaboration outside comments, vote-based validation and basic content modification. The possibility to create and advertise groups of citizens working on a specific topic or living in a given

neighborhood can be seen as a first step towards the acknowledgment of this collaboration by maintainers of participatory platforms. However, as stated in § 4.3.3, the means for collaboration are found outside of online participatory platforms, in online CSCW platforms or in-person meetings [5].

Finally, the objective of co-construction of projects intrinsically creates a risk of information overload, both for contributors and for administrators of participatory platforms [13, 75]. Indeed, the existence of numerous contributions and replies to contributions presents an issue concerning the scalability of participatory processes themselves. This issue is accentuated by the interactive nature of the co-construction of projects, even though it can be also faced when the objectives are different. While Interviewee #2 and Interviewee #3, who work for the smallest platforms of our panel concerning the number of contributions and do not propose co-construction of projects, state that they simply read all the contributions published in their platforms, other interviewees need to use more complex methods to summarize all contributions. Interviewee #1 states that they “do a lot of data and natural language processing using Modalisa³ to help the involved services of the municipality”. Facing the same problem when they work for consultations with high volumes of contributions which have to be processed quickly, Interviewee #4’s company “is working on a R&D project to use contemporary Natural Language Processing techniques on citizen consultations”. Interviewee #5 states that they follow the same goal to reduce this scalability issue using different techniques assessed in [13]:

What we want now is to better take advantage of all this information. Indeed, there is a project [...] for using machine learning techniques to analyze the comments and proposals and other information in the database to extract patterns, trends, sentiment analysis, visualization and present this information to the different units that manage the topics so they can use them easily, instead of ready and manually processing thousands of contributions so that they can define their politics. [...] So far, we have accumulated a lot of data, but is really difficult to manage it and do something with it.

4.4 Different types of platforms for contemporary online citizen participation

During our interviews, we asked interviewees about the functionalities of the platforms they work for. These platforms can be divided in two categories: “suggestion box”-like and interactive platforms. Categorization of the platforms for which our different interviewees work are displayed in Table 4.4.

4.4.1 “Suggestion box”-like platforms

In “suggestion box”-like platforms, the functionalities are designed to facilitate the sending of proposals from the citizens to the organizing entity, in a vertical workflow. These platforms allow little to no interaction between contributors beyond the public sharing of proposals. The content produced by contributors is generally

³Software tool for the analysis of interviews and survey results.

Table 4.4: Categorization of platforms maintained by interviewees according to their features

INTERVIEWEE #	PLATFORM FEATURES
#1	Multiple use cases in the same platform. Participation is implemented both with and without interaction between citizens.
#2	Participation without interaction between citizens inside the platform.
#3	Participation without interaction between citizens.
#4	Multiple platforms both with and without interaction between citizens.
#5	Multiple use cases in the same platform. Participation is implemented both with and without interaction between citizens.

an answer to an open-ended question or a proposal with few details, matching the “idea collection” phase of the generic interaction workflow supporting participatory budgeting presented in [128] (see Figure 2.2). In these platforms, a large part of the synthesis work is done manually by agents of the institution. Interviewee #3 describes the process for their platform as follows:

A first examination is done by the service focused on relations with users and neighborhood councils. Then, the proposal is sent to the most relevant service [of the institution] which will look in detail at potential elements provided and make a complementary examination, i.e. look for additional data to make an estimate of the needed budget and analyze the related problems. [...] Finally, the investigation file is sent to neighborhood councils, municipal officials, etc.

The interactive dimension of these platforms is very limited since contributors cannot interact with one another’s contributions using comments or other text-based contributions (e.g., submitting a revised version of a proposal). The means of interaction tend to be limited to the ability to signify one’s interest in a proposal by following it or voting in favor of it. To justify this choice of a restrictive workflow, Interviewee #3 states that “the objective is mainly to facilitate the use [of the platform] and to let contributors freely present projects”. This choice of workflow is also the most commonly made when the objective of the platform is to report simple issues to a service –Interviewee #1 states that they use this workflow for their “potholes and dog mess signaling tool”. Furthermore, the absence of means for contributors to interact with one another’s contributions can be related with the concern for misconduct of certain contributors. Interviewee #2 supports this hypothesis by stating that their city “didn’t want a platform working like a social network, i.e with the possibility for people to remain anonymous when they submit a proposal or a comment”. This intuition that anonymous comments leads to incivility tends to be confirmed, with some degree of mitigation, by Santana [145]. Moreover, Interviewee #4 indicates that their participatory platforms are explicitly designed to limit interactions between contributors and justify their decision as follows:

Participatory platforms also aim [...] at avoiding the pitfall of forums. When we created our company, the most common means of participation were forums organized by news websites at the end of their articles. Many of these websites removed this feature because their moderation rate was extremely high and as soon as there were interactions, three people were monopolizing the conversation, talking only to each other and, after a maximum of ten messages, insulting each other.

During our interviews, administrators of these platforms mentioned neighborhood councils or similar structures more frequently than other interviewees. Interviewee #2 states that the designers of the platform they maintain explicitly designed it to mimic the workflow of offline neighborhood councils. Interviewee #3 indicates that the proposals submitted by contributors are reviewed by a neighborhood council before their evaluation by the related institution. These connections with offline equivalents of participatory platforms can be one of the reasons why the platforms produced tend to implement little to no means for interaction between citizens, since aforementioned neighborhood councils are constituted of small groups of citizens without any interaction workflow implemented by the city.

4.4.2 Interactive platforms

The platforms we consider as interactive are the ones designed to allow contributors to interact with one another's contributions in a horizontal workflow. Even though it is not mandatory, in the case of the platforms we have studied, they do also implement a vertical workflow to allow the organizing institution to gather and process contributions and publish replies to some of them. Platforms of that kind generally implement – although not being limited to – the three following forms of contributions:

1. User A submits an initial contribution, detailed or vague, to the platform.
2. User B comments on a contribution formerly submitted by User A.
3. User B votes for, against, or neutrally, on A's contribution.

These platforms tend to implement the “idea collection”, “proposals development” and “deliberation and voting” phases of Parra *et al.*'s generic workflow [128]. Consequently, they allow users to have an impact on a project on further steps of its elaboration than “suggestion box”-like platforms. Interviewee #1 states that in their consultation platform, they mainly act to inform citizens of their impact on the project, giving them in certain cases a wide scope of action in the synthesis process:

We are the guarantors of the proper organization of the consultation, i.e. we have to include the citizen in the project as strongly as possible, to make them understand what questions we are asking, to let them know their degree of implication in the final decision, and to ensure that they receive at least some feedback by the end of the consultation.

As stated in § 4.3.4, interactive platforms tend to be used when the objective of the participatory initiative is to co-construct a project. Interviewee #4, whose company produces participatory platforms using different functional modules, confirms this observation by saying that modules allowing more in-depth interaction between contributors (e.g., the possibility to submit a modified version of another user's contribution) are implemented in a platform only “if they want to produce a structured, problematized and contradictory thought process to identify points of consensus and dissensus concerning a given topic”. Interactive features also tend to be implemented in the cases where the expected number of contributors is high,

like national-scale consultations (e.g., the Universal Activity Income consultation⁴ in France) or participatory platforms for large cities (see Table 4.4). This tendency may be explained by the need for a wide enough audience to make use of such features.

Finally, even though interactive platforms tend to implement more means for collaboration than their “suggestion box”-like counterparts, these means are very limited with in comparison with state-of-the-art online collaborative systems. The only example of collaborative features found in the platforms of our panel is the possibility to create contributions derived from another contributor’s own contribution in certain platforms. However, most features for collaborative editing of contributions (e.g., version control, joint writing...) are not implemented, therefore requiring external tools to produce common contributions as mentioned in § 4.3.3.

4.5 How to enrich online participatory platforms

The analysis of interviews of stakeholders of online citizen participation proposed in § 4.3 and § 4.4 allow us to identify two axes to differentiate participatory platforms: the objectives pursued and the presence or absence of features for interaction between citizens. These interviews also allow us to highlight flaws in some of the categories of participatory platforms identified. Based on pre-existing works, we propose potential solutions to address the aforementioned flaws in future participatory platforms. For the sake of clarity, solutions proposed to issues identified in this chapter are summarized in Table 4.5.

4.5.1 Collaborative systems

The lack of means for collaboration in online citizen participation has been highlighted several times in this chapter. This impossibility for groups to collaborate inside participatory platforms does not encourage contributors to produce common, well-developed contributions. Implementing features for collaborative editing of contributions could be a mean to reduce the raw number of contributions by merging similar ones – therefore reducing the information overload previously identified [38, 132]. It could also increase the degree of participation by giving a wider space to deliberation between contributors than on existing platforms where the means of interaction with contributions are simple (vote, comment and, in certain cases, propose modifications).

The first question towards the implementation of collaborative editing in participation is whether the system has to be *synchronous* or *asynchronous*. In the former case, the changes that a user applies to a document are shared with the rest of the contributors as soon as they are performed, generally as character-based updates. Platforms such as Google Docs and EtherPad implement this method. In the latter case, the users decide when they want to share their contributions with –and receive the ones from– the other users. This is the approach implemented by, e.g., Version Control Systems (VCSs) such as Git [164] or Subversion [114], and the massively used wiki engine Wikimedia. We identify the two following means to enrich online citizen participation using synchronous and asynchronous collaborative editing:

⁴<https://www.consultation-rua.gouv.fr/>

- Using asynchronous collaborative editing features, contributors could have the opportunity to alter one another's contributions to produce revised versions and to merge several proposals together. These features are partially implemented in certain platforms (e.g., *République Numérique*⁵ or *Revenu Universel d'Activité*⁶ consultations) but these platforms lack the tools implemented by contemporary asynchronous editing systems (e.g., VCSs, wikis) like the merging or the labeling of different versions of a document. The choice of asynchronous over synchronous editing in that case is motivated by the possibility to easily keep track of revisions and to prevent issues such as vandalism [3, 151] and edit wars [159, 180]. We acknowledge the fact that, in contemporary use cases, asynchronous editing tends to be implemented when contributors have a high level of computer literacy (e.g., collaborative software development, wiki-based information sharing) and remains little used by the general public. To reduce this issue, certain activist communities⁷ and academic courses [142] organize group training sessions to teach the use of platforms of this type, but the use of synchronous editing systems can also be seen as a solution to this issue.
- Participatory platforms could also leverage the specificities of synchronous collaborative editing (e.g., its ease of use or the possibility to do joint writing) to allow groups of users to collaboratively create a common proposal from the ideation phase to its sharing with the rest of contributors. As shown in [123], synchronous editing applications give users new means of collaboration using recognizable and specific patterns of work. However, the direct sharing of contributions to a document can also be the source of inconveniences for users, from the complexity induced by collaboration rules of a group to the difficulty to share contributions when they are yet not finished and do not fully represent the ideas their authors want to share [170]. We therefore recommend to implement synchronous collaboration only in contexts where contributors share the same objective, like an opinion group. To the best of our knowledge, such features are not implemented in contemporary citizen participation platforms.

4.5.2 Natural Language Processing for large-scale processing of contributions

In this chapter, we highlighted the risks induced by large volumes of contributions in citizen participation platforms, especially concerning the difficulty for platform administrators to produce an accurate summary of a high number of contributions (see § 4.3.4). Previous works [38, 132] have also identified the risk of information overload for citizens when the number of contributions published is too high for them to easily identify the most relevant contributions depending on their topic of interests or areas of expertise. These issues should be given particularly high importance in cases where the objective is the co-construction of a project between citizens and an institution and where means for text-based interaction between contributors are implemented.

⁵<https://www.republique-numerique.fr/>

⁶<https://www.consultation-rua.gouv.fr/>

⁷E.g., <https://www.artandfeminism.org/>

Table 4.5: Summary of the issues identified and means proposed in Section 4.5 to solve them.

ISSUE	PROPOSED SOLUTION	SOURCES
Lack of means for collaboration	Asynchronous collaborative editing in cases where approval of new changes to a proposal is needed and where vandalism or edit wars are major concerns.	[3], [114], [151], [159], [164], [180]
	Synchronous collaborative editing in cases where contributors share the same objective, i.e., vandalism and edit wars are minor concerns.	[8], [123], [170], [172]
Information overload in large-scale participation	Automated tag generation using NLP to help filter contributions and recommend the most relevant contributions for a given user.	[13], [26], [156]
	Automated summarization of large sets of contributions by identifying the most relevant contributions and arguments for a given topic.	[13], [92], [125], [149], [150]

The automation of certain tasks currently performed by administrators of participatory platforms or by contributors themselves can be a mean to address the aforementioned risks. We consider that such automation can be done using different Natural Language Processing (NLP) tools. We identify the two following use cases where machine learning can be leveraged to improve large-scale online citizen participation:

- Participatory platforms could use NLP tools to generate tags for the different contributions published. Implementation of such feature would allow contributors to more easily identify the most relevant contributions based on their topics of interest. Platforms administrators would also be able to filter contributions focused on a given topic and therefore summarize them more efficiently. Contribution tagging can also be used to recommend relevant contributions to a user based on topical similarity with a given first contribution (e.g., the one they are currently reading). Dynamic identification of the most relevant tag words can be done using TF-IDF [156] or similar algebraic methods. However, topical clustering of contributions and contributors can also be done using machine learning-based algorithms such as LDA [26] to probabilistically assess the relevance of each contribution for a given topic. The two aforementioned methods have been experimentally implemented by Arana-Catania *et al.* [13] for the Decide Madrid platform. However, while Interviewee #5 states that they use similar NLP features to help platform administrators summarize participatory processes, none of the contributor-oriented features are implemented in the current version of the platform.
- Another use case for NLP in citizen participation platforms would be the summarization of contributions on a given topic. Indeed, given a large set of contributions on a given topic, both contributors could want to know the most frequent arguments or the most relevant contributions. Using cosine similarity of sentence vectors⁸, Arana-Catania *et al.* [13] produced a network of sentences from all the comments of an initial proposal and applied the PageRank [125] algorithm to produce a summary of the most relevant remarks to this proposal, with mixed results. However, other techniques such as argument mining can be leveraged to produce a collection of arguments about a given contribution or

⁸Sentence vector is a version of a sentence encoded by a sentence embedding technique to allow their use by a language understanding model (e.g., SBERT [138]).

topic [92]. Segura-Tinoco *et al.* [149, 150] propose an encouraging step towards the implementation of such technique in citizen participation contexts.

4.6 Threats to validity

One potential threat to the validity of our study is the presence of online participatory initiatives in other countries or using different platforms than those taken into account by our interviews. Indeed, adding interviews from other citizen participation platforms administrators or from active contributors could improve our findings and add new keys to understanding collaboration in citizen participation platforms. However, the large diversity of the platforms managed by our interviewees –in terms of objectives, functionalities and size of the involved group of potential contributors– allows us to describe a wide spectrum of practices. This diversity of profiles is described in Section 4.2 of the chapter. We consider our set of interviews to be diverse enough to describe contemporary online citizen participation platforms.

A second issue is related to the fact that all our interviewees work for French and Spanish citizen participation platforms. Indeed, the addition of platforms used in countries outside of Europe could add new perspectives based on these countries' legal framework or culture concerning participation. However, the examples of participatory platforms that we found outside of Europe were either out of service (e.g., Argentinian *Partido de la Red's* implementation of the DemocracyOS framework) or too close to other platforms already in our panel (e.g., New York City's participatory budgeting platform⁹ based on Decidim).

4.7 Conclusion

This chapter presents an analysis of online citizen participation platforms based on a series of semi-directed interviews with different stakeholders involved in the creation and maintenance of the said platforms. We propose a classification of contemporary participatory platforms based on two axes: the output they aim at producing and their types of features for interaction between contributors. We identify specific issues for the platforms aiming at co-constructing projects with citizens at a large scale and for those acknowledging the possibility for contributors to produce collective contributions. Finally, we propose to leverage different state-of-the-art technologies – namely collaborative editing and natural language processing – to solve these issues and help produce platforms more suitable for both large numbers of contributors and high degrees of participation [15]. Both the highlighting of issues on the different types of participatory platforms and the identification of state-of-the-art methods to solve the said issues are crucial means to help platforms administrators and designers improve online citizen participation. Therefore, in the next chapter, we will assess the relevance of an NLP-based method to improve large-scale citizen participation.

We consider that identifying the objectives of institutions organizing participatory initiatives as well as their motivations for implementing specific features is essential to understand the design – and potential flaws – of contemporary participatory platforms since they have the most direct influence on their creation. However,

⁹<https://www.participate.nyc.gov/>

we also acknowledge the importance of understanding the issues faced by contributors in online participation contexts to improve their involvement in public decision making. Therefore, we argue that both the inclusion of citizens in the governance of participatory platforms and quantitative analysis of the feedback of contributors [63] are promising leads towards the creation of more viable participatory platforms.

Chapter 5

Improving large scale participation using NLI

Contents

5.1	Introduction	85
5.2	Methodology	87
5.2.1	Problem statement	87
5.2.2	Weakly-labeled data for model training	89
5.2.3	Creating a gold dataset for evaluation	89
5.2.4	Automatically labeling long texts	90
5.3	Datasets	91
5.4	Evaluation	95
5.4.1	Fine-tuned models	95
5.4.2	Insights and error analysis	98
5.5	Assessing the integration of an NLI model into a citizen participation platform	100
5.5.1	Identifying idea communities in a consultation	101
5.5.2	Recommending contributions to citizens during their navigation	103
5.6	Interview with stakeholders	105
5.7	Conclusion	107

This chapter of my thesis is based on our paper “Finding Conflicts of Opinion on Citizen Participation Platforms”, written with Oana Balalau, Rafael Angarita and Valérie Issarny. At the time of submission of my Ph.D. thesis, this paper has not been published yet.

5.1 Introduction

In certain instances, citizen participation platforms allow citizens to see and interact with one another’s contributions, generally in an asynchronous collaborative work-

flow (see Chapter 4). This chapter focuses on these latter cases where collaboration-oriented interactions among citizens are made possible. However, even though citizen participation platforms are valuable means to co-construct elaborated projects on a large scale, several studies have reported that users and decision-makers need help at different stages in the participation effort. In [75], the authors identify several challenges, such as the lack of communication between authorities and participants and unclear criteria used by the authorities for selecting among the proposed solutions, or the unwillingness of collaborators to reach a consensus. Another critical challenge for such projects is information overload, affecting decision-makers and users [38, 132].

This work introduces a tool to detect equivalence and contradiction in users' proposals. Our tool allows users to find proposals of interest - whether equivalent or contradictory - and engage with their contributors, while decision-makers can leverage it to create projects composed of *semantically consistent* sets of proposals. Semantic consistency, as opposed to syntactic consistency implemented by consistency management algorithms for collaborative editing (e.g., [8]), is defined as the absence of contradiction between different texts. The contradictions of interest can be of different natures, ranging from ideological contradiction to numeric disagreements [50]. We focus on projects composed of proposals contributed by different users. Then, we consider a project to be semantically consistent if it does not contain any pair of contradictory proposals. Automatically ensuring semantic consistency will help citizens build joint well-developed, complex projects by highlighting the potential contradictions in their constituent parts. Moreover, detecting equivalent or contradictory proposals will help platform maintainers identify idea communities and summarize their claims.

We introduce an approach based on Natural Language Processing (NLP) to identify pairs of conflicting proposals via text classification. More precisely, this objective can be identified as a Natural Language Inference (NLI) task, also referred to in the literature as textual entailment. Parikh *et al.* [126] define NLI as “the problem of determining entailment and contradiction relationships between a premise and a hypothesis”. In our case, we try to predict whether two proposals are entailed, contradictory or neutral. We then define our methodology to identify models enabling the detection of contradictory proposals. We apply our method to French-speaking citizen consultations, i.e., “*République Numérique*”¹ (Digital Republic, RepNum for short in the following) and “*Revenu Universel d'Activité*”² (Universal Activity Income, RUA in the following). We chose these consultations because they received a large number of proposals (respectively 692 and 1760) and allowed contributors to submit comments for or against any given proposal, which are exploited for model training. Moreover, the contributions that were submitted in these consultations are typically larger than pairs of sentences found in most NLI corpora (e.g., SNLI [30], XNLI [45]). Therefore, we consider our work as an important step toward applying NLI tasks to large texts. To validate our approach, we perform a user study comparing textual entailment with textual similarity, a technique proposed for organizing citizen contributions [13, 19]. Finally, we present an interview with Cap Collectif, the company which has created the participatory platform used in the consultations we are analyzing.

¹<https://www.republique-numerique.fr/>

²<https://www.consultation-rua.gouv.fr/>

This chapter is organized as follows. Section 5.2 defines our methodology. Section 5.3 describes how we build the datasets we use to train our models according to the specifics of citizen participation data. Section 5.4 describes the results obtained using our methodology with the datasets previously presented. Section 5.5 assesses the benefits of using such a tool on two online citizen consultations that took place in France, while in Section 5.6 we present an interview with Cap Collectif, the company that created the collaborative platform hosting the two citizen consultations we chose to implement our methodology. Finally, Section 5.7 concludes with perspectives for future work.

The code and datasets used in the production of this chapter is available at https://github.com/WilliamAboucaya/rua_repnum_contradiction_analysis.

5.2 Methodology

This section presents our methodology to improve online citizen participation platforms using NLI on the contributions produced.

5.2.1 Problem statement

In this work, we focus on a particular type of citizen participation platform whose purpose is to gather opinions and potentially reach a consensus with citizens on a given topic. The discussion topic is usually structured in several subtopics, referred to as sections. We suppose that the platform implements the following workflow: *i*) A discussion starts in a section with a proposal. A proposal has a title (i.e., a short description) and a body (i.e., a detailed description of the proposal). *ii*) Each proposal can receive comments; *iii*) Each comment has a stance towards the proposal, in support or contradiction. From now on, we will refer to platforms fulfilling such requirements for citizen participation platforms. This format is also typical of a debate platform, where proposals are claims for which users present pro and against arguments.

Based on previous works [38, 75, 132], we identify the following needs (**N1**, **N2**, **N3**), expressing issues we aim at mitigating, and requirements (**R1**, **R2**), expressing general criteria to meet, for a citizen participation platform:

(**N1**) Reducing information overload:

- User information overload: a citizen cannot find proposals relevant to their points of view.
- Admin information overload: an admin cannot understand all the points brought forward by the citizens.

(**N2**) Obtaining projects composed of semantically consistent sets of proposals. The absence of contradictions is obtained via conflict resolution between different proposals.

(**N3**) Identifying the main points brought forward by citizens. Identifying groups of proposals which are supporting the same ideas is important for offering a complete overview of a consultation.

- (R1) Quality of results of algorithms integrated into the platform. When employed, algorithms should provide good-quality information to users.
- (R2) Ease-of-use of algorithms integrated in the platform.

These needs and requirements are particularly crucial to maintain healthy citizen participation platforms at large scale. To address the needs and requirements of stakeholders of a citizen participation platform, we propose the use of a method that identifies if two proposals are entailed or contradictory. We validate this approach in Section 5.5.

We consider a pair of proposals to be **entailed** if they mutually support one another’s claims while not having incompatible elements. On the other hand, **contradictory** pairs of proposals support incompatible or compatible claims but describe them with incompatible elements (e.g., different numeric values for the budgeting of the same idea). A pair of proposals is labeled **neutral** if it does not fit the criteria for entailed and contradictory labels, i.e., two proposals supporting unrelated claims without incompatible elements in their descriptions.

Problem statement: *Given a pair of proposals (p_1, p_2) from a citizen participation platform, predict its textual entailment classification, i.e. if the pair is entailed, contradictory or neutral.*

We address this problem by introducing a methodology to efficiently train NLI classifiers for citizen participation platforms based on data produced by their contributors. These classifiers must be able to handle contributions from participatory platforms, which are typically longer than the texts provided by NLI datasets.

We position the methodology presented in this section as a tool to facilitate large-scale online citizen participation in contexts where citizens have the possibility to express their different views on specific debates related to a broad topic (e.g., digital transformation of a country, universal income) defined by the organizer of the participatory initiative. Online citizen consultations chosen as examples in the following section of this thesis, as well as other similar cases³, can be used to implement our methodology since citizens are not bound to answer specific questions or vote on pre-existing items. In these types of platforms, contributions often contain very few elements of structured characterization (e.g., tags) to lower the level of computer literacy required to participate. This leads us to focus on the textual content of the different proposals to better identify their relation with one another – entailed, contradictory or neutral. However, our work is less relevant for other types of online citizen participation initiatives such as citizen surveys or the deliberation phase of participatory budgeting [128]. Other works have been proposed to tackle the specific issues faced in these citizen participation contexts. For example, Goel *et al.* [64] propose a voting method for deliberation in participatory budgeting called “knapsack voting”, presented as an alternative to the most commonly implemented restricted voting methods. Their method is particularly relevant for participatory budgeting in which the objective is to choose how to distribute a finite budget between multiple projects, and allows to better represent the distributions expected by each citizen. However, in citizen consultation contexts, applying complex voting rules is much less likely to fix problems already identified such as information

³e.g., <https://discriminations.cap-collectif.com>, <https://decide.madrid.es>

overload since the objective is to help citizens directly express their different views on a topic through text-based contributions – mainly proposals and comments. On the opposite, our work aims at ensuring that the publication of a large number of text-based proposals is not an obstacle to the proper understanding of the content of a participatory platform by citizens and platforms administrators. This problem is much more relevant for online consultation than for deliberation in participatory budgeting, since in the latter case the number of proposals is typically lower than in large-scale citizen consultation – to the best of our knowledge, participatory budgeting is currently only implemented at the local scale – and certain proposals are generally filtered out or grouped together prior to this step.

5.2.2 Weakly-labeled data for model training

To train classifiers for NLI in citizen participation contexts, we need large datasets of NLI-labeled pairs of texts from similar contexts. However, to our knowledge, such datasets do not exist. Existing datasets for NLI tasks (e.g., XNLI [45], SNLI [30]) are generally composed of pairs of one-sentence statements. Therefore, using the content produced by contributors of citizen participation platforms, we create weakly-labeled datasets for the NLI task. The term **weakly-labeled dataset** is used to denote a dataset created based on rules or human feedback, but for which we do not have a strong guarantee of quality as it was not obtained via an annotation procedure, hence some labels might be incorrect. The objective is to train our models on pairs of texts concerning topics similar to those treated in the real case data – here, the proposals in a citizen participation platform – using a similar vocabulary and writing style and containing a comparable number of tokens.

As stated in § 5.2.1, proposals receive comments explicitly identified as for or against them. To create a relevant NLI dataset, we use pairs constituted of a proposal and one of its comments. We use comments supporting the proposal to produce entailed pairs and comments against the proposal to create contradictory pairs. We select a proposal and a comment associated with another proposal for neutral pairs. Hence we associate proposals with comments about unrelated topics. This weakly-labeled dataset can then be used to fine-tune either a language understanding model (e.g., BERT [52]) or a model already trained for the NLI task.

5.2.3 Creating a gold dataset for evaluation

The weakly-labeled dataset is close to the problem stated in § 5.2.1, which is relevant for training NLI models. However, assessing the different classifiers on pairs of proposals is important for a complete evaluation. Therefore, we create a gold dataset composed of pairs of proposals from the consultation manually annotated using the labels entailed, neutral and contradictory. Pairs of proposals were generated using the following method:

1. First, we select a random proposal from the consultation studied.
2. Then, using SBERT’s multilingual models [137, 139], we compute the selected proposal’s cosine similarity with all the other proposals in the consultation. SBERT is trained for semantic textual similarity and allows us to identify proposals with similar topics.

3. Finally, we retrieve five pairs composed of the initial proposal and its five closest neighbors according to cosine similarity.

We note that choosing random pairs of proposals to annotate would give us a highly imbalanced dataset with a large percentage of neutral pairs. We will show in the following that the top- k closest neighbors of a proposal contain entailed, neutral and contradictory proposals, hence highlighting the relevance of models trained for NLI.

5.2.4 Automatically labeling long texts

Humans are able to understand and remember a lot of contextual information and to make informed decisions based on this information. Proposals can be quite long, and contradictions between two proposals can appear only in a handful of sentences. In this work, we experiment with the amount of context that is necessary for a classifier to make a good classification decision. The entire proposal is the maximum amount of context, but the classifier could be overwhelmed. Hence, we also experiment with smaller units of context, such as a sentence or two consecutive sentences. Therefore, we propose three different **labeling strategies** to automatically label a pair of proposals p_1 and p_2 :

Proposal-wise: The model takes as input a pair of proposals and returns a label we define as the prediction. If the proposals are longer than the tokens limit imposed by the model, then the last tokens are not considered. This strategy is identical to the one used during the training of our models on the weakly-labeled datasets.

Sentence-wise: First, we split p_1 and p_2 into sentences. Then, we feed the model with every possible pair containing a sentence from p_1 and p_2 , respectively. Finally, we compute S_e and S_c , respectively, the share of pairs labeled as entailed or contradictory among all the pairs of sentences from p_1 and p_2 . If S_e (resp. S_c) is higher than a threshold t_e (resp. t_c), then the proposals are labeled as entailed (resp. contradictory). Otherwise, neither S_e nor S_c are higher than their respective thresholds, and the proposals are labeled as neutral. Finally, if both S_e and S_c are higher than their respective thresholds, the label is chosen according to the highest value between $S_e - t_e$ and $S_c - t_c$.

Sentence couple-wise: This strategy is similar to the sentence-wise strategy with the following difference: the pairs used to feed the model are constituted of two consecutive sentences from p_1 and p_2 . The objective is to give more context to the classifier than in the sentence-wise strategy. The label for the pair of proposals is then computed exactly the same way as for the sentence-wise strategy.

For sentence-wise and sentence couple-wise strategies, we compute t_e and t_c for each <model, strategy> pair using the weakly-labeled dataset. First, we compute S_e and S_c for each pair in the test split of the weakly-labeled dataset using either the sentence-wise or sentence couple-wise strategy. Then, we use the SciPy [166] implementation of the dual annealing optimization algorithm to find the (t_e, t_c) pair for which the macro average of $F1$ -scores is maximized for our NLI task. The pair of

Table 5.1: Distribution of NLI labels for pairs of texts used for models training

DATASET	NB. OF ENTAILED PAIRS	NB. OF NEUTRAL PAIRS	NB. OF CONTRADICTIONARY PAIRS
RepNum _{WL}	2,953	2,995	3,037
RUA _{WL}	4,954	4,489	4,024
XNLI _{FR}	133,399	133,400	133,403

thresholds computed on the weakly-labeled dataset of a given consultation is used for pairs of proposals from the same consultation.

5.3 Datasets

We apply the methodology described in the previous section to two major online citizen consultations in France: RepNum and RUA.

The **RepNum consultation** has already been presented in Section 3.2, therefore we will only present its most relevant characteristics for this chapter. This consultation aimed at allowing contributors to propose amendments to a bill on how the French state should be organized in a digital era⁴. The main themes of the consultation were the digital transformation of the French administration, open data, online privacy, and access to the internet. Proposals published in this consultation were regrouped in three different sections depending on their topic. The platform allows contributors to interact with one another’s proposals using, among other things, comments labeled as for or against the proposal.

The **RUA consultation** was held between 10/9/2019 and 11/20/2019 on a dedicated participatory platform. The objective of the consultation was to help the French government gather citizens’ opinions on a new solidarity system based on three main axes: (i) the objectives the system should pursue, (ii) its administrative and financial functioning, and (iii) the criteria to benefit from this new income. At the beginning of the consultation, 18 proposals had been submitted by the French Ministry of Health and Solidarity, followed by 1,742 additional proposals submitted by citizens and other stakeholders. Similarly to *RepNum*, this consultation was held on a platform based on the proprietary solution for citizen participation developed by the Cap Collectif company. Proposals published have a title and a body text explaining their content. The three axes of the consultation were split into eight sections regrouping proposals about a specific topic. The consultation took place on a platform based on the same proprietary solution as the RepNum consultation. To our knowledge, no prior academic research work has been published on this consultation. Contributions to the *RUA* consultation and related metadata are available online as open data.⁵

Datasets for model training. To train our models, we used three different datasets, namely the French-speaking section of XNLI (named XNLI_{FR}) and two weakly-labeled datasets we produced from the RepNum and *RUA* consultations (respectively named RepNum_{WL} and RUA_{WL}).

⁴https://en.wikipedia.org/wiki/Loi_pour_une_République_numérique

⁵<https://www.data.gouv.fr/fr/datasets/consultation-vers-un-revenu-universel-dactivite-1>


Création d'une plateforme de modification des textes officiels.
122 votes • 1 amendement • 11 arguments • 0 source

Explication de l'article

Notre démocratie actuelle en est de moins en moins une. Le citoyen ne peut plus s'exprimer que par son vote. Il est temps de lui permettre de donner son avis sur les textes qui régissent sa vie. Cela doit se faire sans intermédiaire, directement à travers une plateforme web.


1° À l'aide d'une plateforme ouverte à tous, le gouvernement doit permettre au citoyen de proposer des modifications sur les textes officiels.
2° Ces textes peuvent être des lois, la constitution ou des codes.
3° Les acteurs de l'État s'engagent moralement à lire et prendre en compte les propositions.
4° Les acteurs sus-cités doivent être présents sur la plateforme pour participer au débat.
5° Cependant il ne doivent pas avoir un accès administrateur, leur voix doit être égale à celle d'un citoyen.
6° Chaque mois une commission composée de citoyens tirés au sort et de membres du gouvernement doit se réunir afin de discuter des propositions.
7° En cas de modification importante du texte, un référendum peut avoir lieu.

6 arguments pour Les plus récents ▾

 18 octobre 2015 à 22:35


7° En cas de modification importantes, un référendum DOIT avoir lieu .

D'accord 0 Signaler Partager ▾

 17 octobre 2015 à 23:04


Ceci est une consultation qui permet d'enrichir le débat parlementaire et non le concurrencer. Cela permet un débat citoyen et sans doute émergent des idées que ne pouvaient pas avoir les députés. De plus, avec les personnes parachutées, la disponibilité des élus est toute relative.

D'accord 2 Signaler Partager ▾

 15 octobre 2015 à 16:42


oui mais seulement si il y a identification correct d'un citoyen français (car aujourd'hui n'importe qui peut s'inscrire même les gens qui ne sont pas de ce pays)

5 arguments contre Les plus récents ▾

 17 octobre 2015 à 15:52


Ce n'est pas au citoyen de faire cela, mais au législateur.

D'accord 1 Signaler Partager ▾

 8 octobre 2015 à 13:20

Une plateforme de ce type, et c'est déjà le cas sur celle-ci, engendre des propositions sans rapports qui deviennent difficiles à étudier.

D'accord 2 Signaler Partager ▾

 28 septembre 2015 à 17:12

"citoyens tirés au sort et membres du gouvernement" ? Non non : il y a le pouvoir législatif pour cela, élu par tous les citoyens, sur le principe : 1 citoyen = 1 voix. Et les électeurs ne sont pas tirés au sort

D'accord 1 Signaler Partager ▾

Figure 5.1: A proposal (top) with the arguments pro (bottom-left) and against (bottom-right) it in the RepNum consultation platform (anonymized).

Translation of the proposal: Our country is becoming less and less democratic. The citizen can only express themselves by his vote. It is time to allow them to give their opinion on the texts which govern their life. This must be done without any intermediary, directly through a web platform. 1° Using a platform open to all, the government must allow the citizen to propose modifications on official texts. 2° These texts can be laws, the constitution or codes. 3° The actors of the State are morally committed to read and take into account the proposals. 4° The above-mentioned actors have to use the platform to participate in the discussion. 5° However, they must not have administrator rights, their voice must be equal to that of a citizen. 6° Every month a committee composed of citizens picked at random and members of the government must meet to discuss the proposals. 7° If a text undergoes significant modifications, a referendum may be conducted.

Translation of the first comment pro: 7° In case of important modifications, a referendum MUST be held.

Translation of the second comment pro: This consultation allows to enrich the parliamentary debate and not to compete with it. It allows a citizen debate and undoubtedly ideas that representatives could not have will appear. Moreover, with parachuted people, the availability of the elected representatives is quite limited.

Translation of the third comment pro: Yes but only if identification of a French citizen is ensured (today anyone can register even people who are not from this country)

Translation of the first comment against: It is not the citizen's job to do this, but the legislator's.

Translation of the second comment against: A platform of this type, and it is already the case on that one, generates irrelevant proposals that are difficult to review.

Translation of the third comment against: "citizens picked at random and members of the government"? No no: it is the job of the legislative power, elected by all the citizens on the principle of "1 citizen = 1 vote". And the voters are not picked at random

XNLI_{FR}: The XNLI dataset [45] is a cross-lingual dataset for NLI tasks. It consists of pairs of sentences originally written in English and translated into 15 languages, including French. Each pair of sentences is annotated as entailed, contradictory or neutral. The French-speaking section of the XNLI dataset contains 400K datapoints. Each datapoint contains a pair of sentences, called the premise and the hypothesis, and a label. The distribution of labels in the XNLI_{FR} is available in Table 5.1. This dataset is by far the largest we have used, but the natural-language texts are drastically smaller than standard texts written in online citizen consultations, with the median number of tokens being 23 for premises and 13 for hypotheses. The XNLI_{FR} dataset is split as follows: a training set ($\approx 98.1\%$ of the datapoints), a validation set ($\approx 0.6\%$), and a test set ($\approx 1.3\%$).

RepNum_{WL} and RUA_{WL}: The RepNum_{WL} and RUA_{WL} datasets are weakly-labeled datasets created from the textual content of the associated consultations. As said earlier, these consultations allow contributors to add comments to a given proposal and to label them as “for” or “against” the respective proposal (see Figure 5.1). The distribution of labels in these datasets is also available in Table 5.1. These datasets are smaller than XNLI_{FR}. RepNum_{WL} contains 8,985 datapoints and RUA_{WL} contains 13,467 datapoints. However, the texts are longer: in RUA_{WL}, the proposals have a median number of tokens of 58 and the comments of 42, while in RepNum_{WL} the proposals have a median number of tokens of 234 and the comments of 40. Both datasets are split as follows: a training set (80%), a validation set (10%), and a test set (10%).

Datasets for model evaluation. Even though the RepNum_{WL} and RUA_{WL} datasets are large enough to be used for the training of our models, they are not representative of our real case problem where we want to identify entailed, contradictory and neutral pairs of proposals. To assess our models on data more closely related to our problem, we manually annotated 239 pairs of proposals from the RUA consultation and 107 pairs from the RepNum consultation, using the three labels, entailed, neutral and contradictory. Table 5.2 displays examples of pairs of proposals for the three different labels. These annotations represent our gold dataset, see Table 5.3 for the distribution of labels.

N.B.: The pairs in Table 5.3 have been produced using the method described in § 5.2.3. The high number of pairs of sentences annotated as neutral allows us to infer that cosine similarity by itself is not a meaningful indicator of the entailment or the contradiction of a pair of proposals.

Both datasets were annotated independently by two authors of the original paper, and agreement among annotators was measured using Cohen’s kappa (κ) coefficient [41]. According to Landis and Koch’s classification of κ score [88], our agreement is *substantial* ($\kappa > 0.6$) for both datasets and therefore, we consider that comparison with our labels is a valid means to evaluate our models. Furthermore, all the datasets are available in two versions: with and without the title of the proposals before their body text, but the annotators have worked on the version with the titles included. We note that the initial pool of sentence pairs contained a small number of contradictory labels. Therefore, a keyword-matching method was used to

Table 5.2: Examples of pairs of proposals from the RUA consultation with their gold label. Proposals translated by this work’s authors. Opinions expressed do not reflect those of the authors of this work.

FIRST PROPOSAL	SECOND PROPOSAL	LABEL
<p>Information and accessibility. Consumer information videos must be accessible to everyone, including deaf people. This means ensuring the presence of subtitles and French sign language interpretation. All consumer information posted by the relevant authorities and related to any government decision must be fully accessible.</p>	<p>Practices - Accessibility. Implement subtitles or French sign language translation for deaf and hearing-impaired consumers to better inform them concerning company policies and consumers’ rights.</p>	Entailed
<p>An evaluation by a social worker (based half on an evaluation grid, half on human investigation) and multi-disciplinary support. A person must be assessed by a well-trained social worker to give them right to the Universal Activity Income and multi-disciplinary support must be offered according to their needs: support towards employment, help with budget management, support towards care or adaptation of housing (if unsafe or inaccessible). Working in partnership with the different stakeholders would make it possible to follow the evolution of the beneficiaries without putting the burden of monitoring on a single worker. It would be much more beneficial and appropriate.</p>	<p>Automated payment. The payment of the Universal Activity Income should not be subject to a recurrent resources/status reporting. Interconnections are possible (or even already exist between tax collection services, family allowance funds, social security, employment agencies, etc.) so that the calculation of the amount and the payment are made automatically. This would also avoid overpayments due to errors or delays in payment that could place the poorest people in even more difficult situations (eviction from housing, bankruptcies, etc.).</p>	Contradictory
<p>Taking into account people who live in couples without being married for reversionary pension! Hello! it seems normal to me that people who live in couple without being married for years and who have 1 or more children can benefit from the reversionary pension of the spouse! Nowadays marriage does not mean anything anymore, it is just an administrative act!</p>	<p>Take into account the income of each person rather than of the couple. Today, for certain aids, the income taken into account is that of the household, therefore reducing the aids for couples. If we use individual income, each person will be entitled to what is due to them according to their situation. The "adult with a disability" allowance for a couple is not the addition of two allowances for single people. Moreover, at age 60 (automatically retired), their income decreases but the disability is still there...</p>	Neutral

Table 5.3: Distribution of NLI labels in the gold dataset

CONSULTATION	NB. OF ENTAILED PAIRS	NB. OF NEUTRAL PAIRS	NB. OF CONTRADICTIONARY PAIRS	κ AGREEMENT
RepNum	22	65	20	0.614
RUA	94	108	37	0.621

find additional contradictory examples. As a result, twenty additional contradictory pairs have been added to RUA, while twelve contradictory pairs have been added to RepNum. The proposals (body and title) in the manually annotated dataset for the RepNum (resp. *RUA*) consultation have a median length of 179 tokens (resp. 87 tokens).

5.4 Evaluation

5.4.1 Fine-tuned models

CamemBERT [105] is a state-of-the-art language model for French. Moreover, unlike certain other models such as GPT-3 [31], CamemBERT is available under an open source license, which is considered as a crucial feature for trustworthiness of participatory platforms [25]. CamemBERT has been pretrained on masked language prediction and is fine-tuned on the required downstream application. In this work, we propose several fine-tuning scenarios, and we investigate which has the best performance. We use CamemBERT_{LARGE}, a neural network with an architecture of 24 layers, 1024 hidden dimensions, 12 attention heads, and 335M parameters. We fine-tune the CamemBERT model using the following seven combinations:

RepNum: CamemBERT model fine-tuned on the RepNum_{WL} dataset.

RUA: CamemBERT model fine-tuned on the RUA_{WL} dataset.

RepNum + RUA: CamemBERT model fine-tuned on both the RepNum_{WL} and RUA_{WL} datasets.

XNLI: CamemBERT model fine-tuned on the XNLI_{FR} dataset.

XNLI → RepNum: CamemBERT model fine-tuned on the XNLI_{FR} dataset, then fine-tuned again on the RepNum_{WL} dataset.

XNLI → RUA: CamemBERT model fine-tuned on the XNLI_{FR} dataset, then fine-tuned again on the RUA_{WL} dataset.

XNLI → RepNum + RUA: CamemBERT model fine-tuned on the XNLI_{FR} dataset, then fine-tuned again on both the RepNum_{WL} and RUA_{WL} datasets.

We train our models for 5 epochs, with a learning rate of $1e-5$ as suggested in [105]⁶, a weight decay of 0.01 and a batch size of 8. We implemented our models using PyTorch [129] and the Transformers library [177].

⁶Default parameter proposed for the fine-tuning of RoBERTa on the MNLI task (see <https://github.com/pytorch/fairseq/blob/master/examples/roberta/README.glue.md>).

Table 5.4: Results on the RUA gold dataset: macro F1-score and precision on entailed label (P_E) and contradictory label (P_C) for the manually annotated pairs of proposals of the RUA consultation. Results obtained using only the content of the proposals, except starred (*) results where the title of the proposals and name of their sections were included. The macro F1-score of a random classifier is 31.01.

MODEL TRAINED ON...	LABELING STRATEGY																	
	WITH PAST TENSE SENTENCES									WITHOUT PAST TENSE SENTENCES								
	PROPOSAL			SENTENCE			SENTENCE COUPLE			PROPOSAL			SENTENCE			SENTENCE COUPLE		
	F1	P_E	P_C	F1	P_E	P_C	F1	P_E	P_C	F1	P_E	P_C	F1	P_E	P_C	F1	P_E	P_C
RepNum	30.8	43.1	29.0	31.8	44.8	20.4	37.3	49.3	17.6	30.8	43.9	27.8	35.8	41.8	26.9	37.9	48.8	20.0
RUA	45.7	55.0	29.6	43.5	53.9	55.6*	35.6	50.0	22.2	44.2	52.8	30.8	43.1	52.9	55.0*	35.7	51.7	28.6
RepNum + RUA	44.8	50.9	28.9	50.0	59.5	26.2	36.6	55.6	20.0	45.2	51.0	32.5	49.5	58.7	28.2	36.8	52.1	18.2
XNLI	45.7	76.5*	51.9	45.1	58.3	32.9	37.9	76.2	27.0	45.1	74.2*	50.0	42.7	54.8	32.9	37.9	76.2	27.8
XNLI → RepNum	30.6	44.4	20.5	39.1	46.9	23.7	40.6	48.4	31.8	30.3	45.2	20.5	39.5	45.2	25.6	39.9	47.8	31.8
XNLI → RUA	52.4	58.9	35.7	54.7	59.3	39.4	38.8	54.3	26.7	53.7	57.8	43.2	54.7	59.3	42.4	39.4	57.1	27.6
XNLI → RepNum + RUA	48.2	55.9	32.5	51.1	61.3	29.8	40.6	55.0	35.7	48.8	55.0	40.0	48.9	59.1	28.9	41.7	55.2	38.9

Once the models have been trained, we apply the strategies described in Section 5.2.4 to two different versions of the datasets: one containing the complete proposals without any pre-processing and another where past tense sentences were removed. This pre-processing is motivated by the fact that in French, sentences expressing instructions – i.e., the components and objectives of the law, that is, the elements of potential disagreement with other proposals – are written in present or future tense. Consequently, past tense sentences in our datasets can be seen as elements of context and, potentially, as superfluous and unnecessarily increasing the number of tokens. To perform this pre-processing, we trained a token classifier for Part-Of-Speech (POS) tagging using the PERCEO corpus [17], where verbs are labeled using the “VER” tag and a sub-tag identifying their tense. Then, we apply our POS-tagger to each sentence of our dataset, and we remove the sentences containing *passé simple*, *imparfait*, and *subjonctif imparfait* verbs since we identify them as past tense sentences.

The results obtained using the 42 different combinations of our fine-tuned models (7 models) and strategies on our two manually annotated datasets (6 strategies) are available in Tables 5.4 and 5.5. We compute the macro F1-scores for each combination and the precision for the labels entailed (P_E) and contradictory (P_C). The first metric informs us of the overall performances of a <model, strategy> pair. The second and third metrics allow us to assess the model’s suitability for the large-scale evaluation presented in § 5.5. We analyze the performance of the two datasets separately.

Given the definition of a contradictory pair we use (see § 5.2.1), one could argue that the sentence-wise and sentence couple-wise should label a pair of proposals (p_1, p_2) as contradictory if at least one of the sentences of p_1 is contradictory with any sentence from p_2 , i.e., if $S_c > 0$. However, given the consistently low results for (P_C) obtained using this method, we have chosen not to include it in our experimental evaluation. For example, using the RUA model and sentence-wise strategy on the

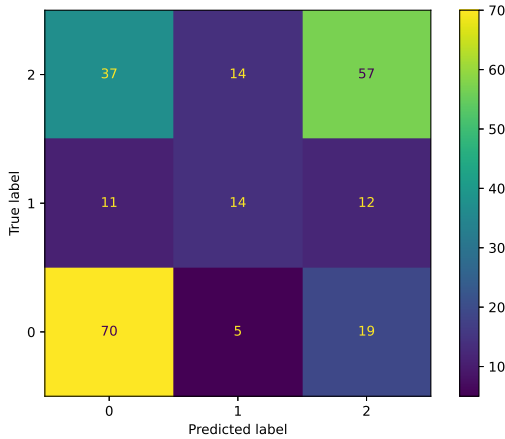


Figure 5.2: Confusion matrix for the RUA manually annotated dataset without past tense sentences, using the “XNLI \rightarrow RUA” model and sentence-wise strategy. The labels are as follows: entailed (0), contradiction (1), neutral (2).




Figure 5.3: Confusion matrix for the RepNum manually annotated dataset without past tense sentences, using the “RepNum + RUA” model and proposal-wise strategy. The labels are as follows: entailed (0), contradiction (1), neutral (2).

RUA gold dataset, $P_C = 15.7$, which is significantly lower than the results obtained with other strategies.

Results for RUA The best-performing model on RUA, according to the $F1$ score on our gold dataset, is the model fine-tuned first on XNLI and then on RUA_{WL} . For each labeling strategy, the results obtained using this model are consistently among the best, if not the best. We observe that fine-tuning on RUA_{WL} generally improves the classifier’s performance, which is expected as this dataset is the most similar to the gold dataset. Using this model and the sentence-wise strategy, independently of the inclusion of past-tense sentences, we obtain a macro $F1$ -score of 54.7. Furthermore, the confusion matrix (see Figure 5.2), which displays the distribution of datapoints depending on their true and predicted labels, shows that this combination of model and strategy struggles with identifying contradictory (1) pairs and on the opposite, is very effective in identifying entailed (0) pairs. However, we cannot identify a labeling strategy that consistently produces the best results for every model. Therefore, for every model, all the strategies (see § 5.2.4) should be tested to find the one giving the best results.

Results for RepNum The best-performing model on RepNum according to the $F1$ score on our gold dataset is the model fine-tuned using both the $RepNum_{WL}$ and RUA_{WL} datasets. The best labeling strategy for this dataset is proposal-wise with past tense sentences removed, with the macro $F1$ -score of 48.8.

The confusion matrix (see Figure 5.3) shows that neutral (2) and contradictory (1) pairs tend to be falsely identified as entailed pairs, while entailed (0) pairs tend to be correctly identified. Overall, the best results are mostly obtained using the proposal-wise strategy. This is an interesting finding, as the proposals from the RepNum consultation are on average longer than those from the RUA consultation (see Section 5.3). We hypothesize that this happens because of how the thresholds

Table 5.5: Results on the RepNum gold dataset: macro F1-score and precision on entailed label (P_E) and contradictory label (P_C). The results are obtained using the content of the proposals prepended with their titles. The macro F1-score of a random classifier is 26.3.

MODEL TRAINED ON...	LABELING STRATEGY																	
	WITH PAST TENSE SENTENCES						WITHOUT PAST TENSE SENTENCES											
	PROPOSAL			SENTENCE			SENTENCE COUPLE			PROPOSAL			SENTENCE			SENTENCE COUPLE		
	F1	P_E	P_C	F1	P_E	P_C	F1	P_E	P_C	F1	P_E	P_C	F1	P_E	P_C	F1	P_E	P_C
RepNum	35.6	27.8	20.0	26.1	23.8	25.0	23.1	22.3	33.3	34.3	26.7	20.0	30.6	26.3	30.8	27.6	22.8	50.0
RUA	42.6	25.4	83.3	30.2	24.4	25.0	33.4	27.0	23.1	40.8	25.0	80.0	30.8	23.7	25.0	33.2	24.4	37.5
RepNum + RUA	47.2	29.0	54.5	37.0	26.0	40.0	45.1	30.8	42.9	48.8	30.0	53.8	33.9	24.4	44.4	43.3	29.0	44.4
XNLI	35.8	22.6	33.3	27.9	20.4	18.2	38.0	32.1	23.8	36.4	25.0	33.3	27.9	19.2	21.7	38.2	34.6	23.3
XNLI → RepNum	45.7	28.6	50.0	31.5	22.0	60.0	31.5	24.4	28.6	45.1	26.9	62.5	31.3	21.2	75.0	32.9	24.1	40.0
XNLI → RUA	39.5	24.1	61.5	32.6	22.5	41.7	35.7	25.8	26.9	42.0	25.0	64.3	25.3	20.2	38.5	37.0	26.7	28.6
XNLI → RepNum + RUA	45.7	30.4	41.2	29.8	21.7	66.7	38.5	23.5	75.0	44.3	29.8	42.9	32.1	21.7	50.0	43.1	25.3	75.0

on the percentage of entailed/contradictory pairs are computed. We recall that we chose the threshold using the weakly-labeled dataset to avoid biased results by computing it on the test set. However, in the weakly-labeled dataset, we have comments which are much shorter than proposals on RepNum. Hence, we might not be choosing the correct thresholds. This hypothesis is confirmed by the fact that, simply by changing the value of one of the two thresholds, we did reach significantly higher values for the macro F1-score – e.g., 58 for the XNLI → RepNum + RUA model with the sentence couple-wise strategy, while the value obtained with the pair of thresholds computed is 38.5.

5.4.2 Insights and error analysis

The results of our different <model, strategy> pairs allow us to draw two conclusions concerning NLI for citizen participation. First, as observed on the RepNum gold dataset, a model can learn using data from another consultation. Indeed, the best results for the metrics computed in this chapter – macro F1-score, precision on entailed, and contradictory labels – on the RepNum dataset are obtained using the RepNum + RUA, XNLI, and RUA models, respectively. Therefore, we consider the datasets produced for this chapter to be relevant as an initial corpus for future works on NLI for collaboration in the context of citizen participation (e.g., in the context of a platform where no weakly-labeled data can be extracted).

Second, confusion matrices highlight that classifiers tend to falsely label neutral pairs as entailed pairs, while the opposite is uncommon. This issue highlights a low recall on the neutral class for most classifiers. This could be because entailed pairs are selected in the weakly labeled data as pairs of proposals and comments “for”. However, a user might add a comment “for” to express an opinion that is on topic, but not addressing the main claims of the proposal, hence being according to our definition a neutral statement. These observations are validated by the fact that a model has a higher recall when trained on XNLI - a dataset where annotators

Table 5.6: Examples of pairs of proposals incorrectly labeled by the “XNLI \rightarrow RUA” model with sentence-wise strategy. Proposals translated by the authors of this work.

FIRST PROPOSAL	SECOND PROPOSAL	TRUE LABEL	PREDICTED LABEL
<p>Review the person’s situation regularly.</p> <p>A person’s Universal Activity Income should be updated regularly to adapt to changes in their family or personal situation. It is also necessary to support people who go back to work without immediately taking away this income. When you get a new job, you have additional expenses and you need time to organize yourself, especially when you have children.</p>	<p>Allow the vulnerable self-employed entrepreneurs to have a safety net.</p> <p>The Universal Activity Income should encourage people to have a salaried or non-salaried activity. From my experience, it is hard to have a non-salaried activity and to receive aid for the difficult months. It is not normal to discourage people who want to start their own business.</p>	Neutral	Entailed
<p>Everyone is equal so the financial aid should be the same.</p> <p>800 € paid automatically to the persons living in France and older than 16 on the 1st of each month without any condition, no fraud possible, no paper to provide.</p>	<p>Universal income for all from birth.</p> <p>Universal income from birth for all citizens legally residing in France; for a child, the payment will be blocked and made available upon reaching legal majority.</p>	Contradictory	Entailed

gave entailed and neutral pairs. For example, for the RepNum dataset, the average recall on the neutral label for the \langle model, strategy \rangle pairs using the XNLI model is 41.2, while it is 30.7 for the pairs using another model. Therefore, new methods to produce better quality weakly-labeled neutral pairs can be beneficial in the training of models.

False identification of neutral pairs as entailed is frequent in cases where both proposals have common objectives but different contexts. For example, the first row of Table 5.6 displays two proposals that aim to mitigate financial risks induced by changing life situations using a temporary income. However, these two proposals describe two different categories of persons to protect- persons who are in charge of a family and self-employed workers- which leads us to label the pair as neutral. This issue is common and is the first lead to improving the reliability of our classifiers to label neutral pairs. Another frequent issue in labeling concerns contradictory pairs of proposals falsely labeled as entailed. An example is shown in the second line of Table 5.6. Here, while the two proposals strongly agree on the principle of a universal income with as little restriction as possible, they differ on the minimum age required to receive the income – respectively starting at the age of 16 and at birth. This difference leads us to label the pair as contradictory while their agreement on the principle of an unconditional universal income probably leads the classifier to

label them as entailed. These examples can be considered as “contextual errors”, i.e. examples where both proposals advocate for similar measures but within different contexts or conditions. We, therefore, consider that higher quality of understanding of the context would improve the quality of models for NLI on texts larger than one sentence.

5.5 Assessing the integration of an NLI model into a citizen participation platform

In this section, we perform a large-scale investigation of the potential of our work. In particular, we will address how we can deal with the **requirements and needs** that we defined in Section 5.2.1. The proposed models can be used to identify groups of proposals defending similar ideas or proposals that are in contradiction, hence they can help solve N1 and N3. Moreover, finding entailed proposals can help citizens discover like-minded contributors and produce joint contributions, while platform maintainers can identify opinion groups to facilitate the synthesis of a consultation. In addition, retrieving contradictory proposals is the first step toward resolving the conflict and solving N2.

The thresholds for the percentage of entailed and contradictory pairs (S_E and S_C) for the sentence-wise and sentence couple-wise strategies are computed in Section 5.4 such that we maximize the macro F1 score. However, one can optimize the thresholds for improved precision to show users only the pairs of proposals labeled with the highest certainty, that is, the pairs of proposals that are more likely to be contradictory or entailed, depending on the purpose. In addition, we can choose a model based on its precision on a specific class, not the overall F1 score.

For the RUA consultation, the best classifier for the class entailed is the XNLI model and the proposal-wise strategy. For the RepNum consultation, none of the classifiers produced have a high enough precision (we aim for a precision higher than 60%) for the entailed label to be reliable for this use case. Therefore, we use a modified version of the classifier composed of the “XNLI \rightarrow RepNum + RUA” model and the sentence couple-wise strategy. We increase the t_e threshold, which improves the precision ($P_E = 66.7$) while reducing the recall ($R_E = 28.6$) for the entailed label. We follow the same approach to select the best models for finding contradictory proposals. For the RepNum consultation, we use the RUA model and the proposal-wise strategy. For the RUA consultation, the classifiers need a higher precision for the contradictory label to be used for contradictory pairs detection. Consequently, we slightly increase the t_c threshold, improving the precision ($P_C = 81.8$) and slightly reducing the recall ($R_C = 24.3$) for the contradictory label. We note that the contradictory model has a very high precision, meaning that a user who asks to see a proposal that contradicts a given one will more likely see a correct answer. For the entailed class, the precision is lower, however, we have shown in our error analysis (Section 5.4.2) that the proposals are still relevant to each other. Hence, we believe that this work is in line with requirement R1.

Once the classifiers have been chosen, we split the consultations using the different sections defined by the participatory platforms (see § 5.3). We then apply the classifiers to all the possible pairs of proposals inside the same section. Results for both the RepNum and RUA consultations are available in Table 5.7. For RepNum,

Table 5.7: Distribution of NLI labels for pairs of proposals in the RepNum and RUA consultations. Overlapping proposals are the ones labeled as entailed by the classifier chosen to maximize P_E and labeled as contradictory by the classifier chosen to maximize P_C .


CONSULTATION	TOTAL NB. OF PAIRS	ENTAILED PAIRS		CONTRADICTIONARY PAIRS		OVERLAPPING PAIRS	
		NB.	%	NB.	%	NB.	%
RepNum	86,398	210	0.2 %	1,548	1.7 %	5	0.006 %
RUA	295,796	28,582	9.6 %	9,659	3.2 %	736	0.2 %

we observe a relatively small percentage of entailed pairs compared to contradictory pairs. RepNum has a total of 692 proposals compared to RUA, which has 1,760 proposals. The small percentage of entailed pairs indicates that RepNum users have more easily identified proposals that correspond with their view on the platform. Consequently, they concentrate their effort on adding contradictory proposals when they disagree with the existing ones. On RUA, we observe a large percentage of entailed pairs, which could result from the difficulty in navigating the larger number of proposals. Finally, we note that many proposals are entailed or contradictory on both platforms, highlighting the important contribution that an NLI model would bring to a citizen participation platform.

5.5.1 Identifying idea communities in a consultation

A first use case for our models would be identifying communities of contributions defending similar ideas. This use case would be relevant to draw an initial summary of the ideas evoked in a given participatory platform when exhaustive reading of contributions is not an option (e.g., when the number of contributions produced is too high). Indeed, in this situation, the automated identification of communities would allow platform administrators to read only a few contributions in each community to identify the main ideas defended.

In the following, we focus on the largest sections regarding the number of proposals for RepNum and RUA. For RepNum, the section title is *La circulation des données et du savoir* (The flow of data and knowledge). For RUA, the section title is *À quels publics le revenu universel d'activité doit-il s'adresser ?* (Who should the universal activity income be aimed at?). We create a graph G for each section where a vertex is a proposal, and two vertices v_1, v_2 are connected if the corresponding proposals were labeled as entailed. Finally, we find communities using scikit-network's [29] implementation of the Louvain algorithm. Proposals belonging to the same community are considered to defend similar ideas. For RepNum, we apply the Louvain algorithm with $\gamma = 1$, and for RUA with $\gamma = 1.5$. We then annotate the proposals in each community depending on their topic and claims, and we compute the purity [102] of the clusters. For RepNum, we also compute the purity of a baseline clustering produced for comparison using the subsections of the section of the consultation – this comparison cannot be performed on RUA as the section studied does not have subsections. Results are available in Table 5.8. Concerning the section from the RepNum consultation, the modularity obtained is high, with well-defined clusters (see Figure 5.4). The section from RUA consultation has a low modularity score, however, both the RepNum and RUA sections have a significant



(a) The nodes are colored and grouped according to the communities found by the Louvain algorithm.

(b) The nodes are colored according to the subsections of the proposals.

Figure 5.4: RepNum proposals from the section “La circulation des données et du savoir”, on the flow of data and knowledge. A node represents a proposal and an edge connects two nodes when the corresponding proposals are entailed. Isolated vertices are not displayed. We observe that entailed proposals are likely to be placed by users in different subsections. A platform implementing our model would be able to guide users towards existing entailed proposals and help them better contribute to the consultation.

level of cluster purity. In particular, the purity of our clustering on RepNum has a very high purity when compared with the baseline clustering. This suggests that the communities in both RUA and RepNum have clearly defined topics and claims widely shared inside the clusters. In Figure 5.4, we observe the communities in RepNum. On the left, we have the distribution of proposals according to the community detection algorithm, while on the right, we see their distribution according to the subsection (the section we selected had 5 subsections). The entailed clusters contain proposals written in different subsections, showing that users might struggle to add contributions in the correct section/subsection, which is coherent with our observations on clustering purity. Hence a platform implementing an NLI model could guide a user toward the best section to make a contribution.

Overlapping pairs of proposals (see Table 5.7), similarly to those labeled as contradictory, are likely to contain two proposals placed in different clusters. Indeed, in the RepNum section, 2 of the 3 overlapping pairs are composed of two proposals placed in different clusters. For the RUA section, the community detection algorithm identifies the proposals of 156 of the 289 overlapping pairs (54.0 %) as belonging to two different clusters. Finding all the contradictory pairs of proposals between clusters of entailed proposals can significantly help decision-makers resolve

Table 5.8: Results of the clustering of the largest section of the RepNum and RUA consultations using the Louvain algorithm.

CONSULTATION	NB. CONNECTED COMPONENTS	NB. OF CLUSTERS	MODULARITY (%)	PURITY (%)	PURITY OF SECTIONS (%)
RepNum	6	13	67.0 %	68.3 %	31.7 %
RUA	1	19	8.5 %	56.5 %	N/A

all potential conflicts on a citizen collaboration platform.

5.5.2 Recommending contributions to citizens during their navigation

Another use case for our models is the recommendation of proposals to citizens to contradict or complement the views of the proposal they are currently reading. Indeed, in the case of large-scale citizen participation, information overload may lead to a higher difficulty for citizens to find different points of view on a given issue. We assess the suitability of our models using the following protocol:

1. We pick 100 random proposals in the RUA dataset. The choice of this specific dataset is motivated by the fact that most citizens have at least some degree of understanding of its topics (e.g., social welfare, taxes), while the RepNum consultation is focused on topics where a large share of the population may need an explanation to understand certain proposals (e.g., data protection, open science). Moreover, we filter our dataset to include only proposals between 30 and 250 tokens for ease of understanding.
2. For each proposal p_i , we pick both the proposal with the highest cosine similarity using SBERT \tilde{p}_i and the contradictory proposal with the highest share of contradictory pairs \bar{p}_i . Here, we use \tilde{p}_i as the baseline result, with which we compare the result obtained using our model. We choose the proposal with the highest cosine similarity rather than the lowest since, while high similarity does not seem to be correlated with entailment (as stated in § 5.3), the proposal with the lowest similarity is almost always off-topic. In cases where \tilde{p}_i was a copy of p_i (e.g., a user posting their contribution in two different consultation sections), we chose the proposal with the second highest similarity score to avoid easily fixed issues.
3. Using Amazon MTurk⁷, we present to each volunteer one of our 100 randomly picked proposals p_i followed by \bar{p}_i . We ask them the following closed-ended question “Does the second proposal (\bar{p}_i) bring new points of view or elements of contradiction to the first proposal (p_i)?” and if their answer is yes, we ask them to list the said new elements brought. This question aims at identifying proposals that are focused on the same topic as p_i while not being paraphrases.
4. We repeat step 3, replacing the second proposal – i.e., \bar{p}_i – with \tilde{p}_i . Steps 3 and 4 are performed in random order to prevent risks of order-related biases. At the end of the experience, the annotator has been confronted to both pairs $\langle p_i, \bar{p}_i \rangle$ and $\langle p_i, \tilde{p}_i \rangle$.

Table 5.9: Results of the MTurk-based experiment assessing the suitability of our model and cosine similarity (baseline) to recommend new proposals to contributors.

TYPE OF PAIRS	TOTAL NUMBER OF PAIRS	PAIRS WITH DIFFERENT VIEWS EXPRESSED	AVG. # OF NEW ELEMENTS IDENTIFIED (ALL PAIRS)	AVG. # OF NEW ELEMENTS IDENTIFIED (PAIRS WITH DIFFERENT VIEWS EXPRESSED)
Contradictory pairs (from our model)	100	76	0.79	1.118
Similar pairs (baseline)	100	49	0.56	1.084

This protocol has been created to replicate the case of a citizen clicking on a proposal p_i and who wants to read other proposals on the same topic – here, \tilde{p}_i and \bar{p}_i – to form their opinion on this topic. We do not include any question about the quality of \tilde{p}_i and \bar{p}_i since it would mainly measure the political stance of the annotators on the different topics presented. To evaluate the quality of \tilde{p}_i and \bar{p}_i independently, a distinct experimental protocol would be necessary. Hence, our evaluation primarily centers on assessing the ability of these two proposals to introduce new perspectives or elements of contradiction to readers of p_i . Each experiment is run with three different annotators to improve the reliability of results. We consider a pair of proposals to express different views if at least two annotators have identified differences of views between them. We compare results obtained using our model to those obtained using cosine similarity since this method has already been evaluated in previous papers for textual document recommendation [175]⁸ and other use cases such as documents clustering [13, 73]. A total of 24 different workers did all the assignments. Results of the experiment are available in Table 5.9. The first and most clear result is that proposals obtained using our model are 55.1 % more often labeled by workers as bringing new elements to the initial proposals than proposals obtained using cosine similarity.

This result leads volunteers to identify 41.1 % more differences of views in pairs of contradictory proposals than in similar proposals. This result implies that recommending proposals obtained using our model –rather than cosine similarity– to contributors will often lead them to read more diverse views on the topics of the proposal they were initially reading, therefore helping them build an informed opinion. We, therefore, assert that our model is more suitable than cosine similarity to produce proposal recommendations in citizen participation platforms. The second result is related to the number of differences of views identified by volunteers. Indeed, the average number of differences of views identified in pairs where different views have been detected is 3.1 % higher for pairs built using our model than for pairs built using cosine similarity. Nevertheless, in this case, the number of pairs for cosine similarity becomes too low to be representative. A potential explanation for this similar number of new views identified may be the fact that proposals used in our experiment have been filtered to contain up to 250 tokens, thus filtering out

⁷The reliability of MTurk for French-language tasks has been evaluated in [131].

⁸Here, we prefer cosine similarity to other similarity measures such as BM25 since it tends to be more reliable on texts larger than one sentence.

the most complex proposals presenting multiple ideas. However, as for now, while our results highlight a slight tendency for pairs produced using our model to contain more differences of views than pairs produced using cosine similarity, we do not draw any conclusion about the ability of our model to identify proposals bringing quantitatively more new views than the baseline model.

These results allow us to infer that our models can be reliably used in citizen participation platforms to recommend proposals to contributors based on the proposals they are currently reading. Indeed, even though the capability of our model to identify proposals with more new elements than cosine similarity is still uncertain, our results clearly show that proposals recommended using our model are more likely than those recommended using cosine similarity to add new elements of discussion to the initial proposal without being off-topic.

5.6 Interview with stakeholders

To assess the relevance of our work and of our use cases for online citizen participation, we conducted an interview with Cap Collectif, the company which has created the participatory platform used in the *RepNum* and *RUA* consultations. This company develops a modular platform to answer the needs of multiple citizen participation use cases (e.g., consultation, participatory budgeting, etc.). Their platform has been used by hundreds of organizations, companies and public institutions, primarily in France, but also across other countries like Canada. Cap Collectif is a stakeholder of online citizen participation, as they represent the administrators of the platform that need to analyse large amounts of textual input from users. Oana Balalau, Rafael Angarita and myself attended this in-person meeting on May 23rd, 2023, with the CEO and the counseling director of the company. The meeting began with a presentation of the works presented in this chapter, followed by a question-and-answer session. Quotes from this interview have been translated from French by myself.

The interviewees have presented the method they currently use to produce the synthesis of an online citizen consultation. First, for each initial proposal submitted by the organizing institution, workers read all the comments posted to produce a more comprehensive list of pro / con arguments and add extra categories to represent more mixed viewpoints, alternative solutions, implementation conditions or points of vigilance. This extensive synthesis of the contributions related to these initial proposals is due to the fact that, according to one of our interviewees, “[their] customers are mainly looking for assessments and arguments related to their proposals”. Then, proposals submitted by citizens in each category of the platform are read and grouped by theme. For each theme, proposals presenting similar ideas are replaced by a short summary of the said ideas. Comments on citizen proposals are dismissed. This tree representation of proposals aims at facilitating the understanding of the different ideas discussed by citizens on the platform.

We have asked our interviewees to quantify the amount of work needed to produce a synthesis of a consultation on a scale similar to the *RepNum* or *RUA* platforms. They gave us the example of the consultation they were currently working on and which had received 1,076 proposals, i.e., a number comparable to the two consultations studied in this paper. For this consultation, they expected the synthesis process to require “two persons working almost full-time for three weeks”. This high

workload induced by fully hand-made synthesis tends to validate our assumption concerning the need for a pre-processing to improve scalability of online citizen participation with regard to the number of contributions. Moreover, at another point in the interview, our interviewees state that they have started implementing NLP for certain tasks. Indeed, to help institutions explore the contributions produced in participatory platforms, they have started to implement sentiment analysis and vocabulary-based contributions clustering. They motivate this decision by stating that “[even though] in general, human work is qualitatively superior [to automated processing], there is a scalability issue so [they] are forced to go in this direction anyway”. We therefore consider that the use of our classifiers to identify idea communities – as proposed in Section 5.5.1 – is a solution to real issues faced by participatory platforms administrators and participatory processes organizers and can be integrated to the synthesis work.

Interviewees consider that the use of our tool for this use case is also relevant in other contexts than citizen consultation. They propose the example of “suggestion box”⁹ platforms in which, at the time of the interview, they were trying to implement proposals clustering with similar objectives to those presented in Section 5.5.1. Moreover, according to our interviewees, “once 50 % of the contributions have been read, sometimes even 20 %, new themes are no longer really created, contributions are simply added to the existing ones”. This leads us to propose the following example protocol leveraging idea communities identified by our classifiers in an online participatory platform to reduce the workload faced by people in charge of the synthesis work. For each cluster identified, we propose to only read an arbitrary percentage of the proposals chosen randomly. Our clustering allows us to reduce the risk to lose certain views because of the random sampling. A possibility to further reduce this risk would be to lower or increase the percentage of proposals read on a specific cluster based on the purity of the first proposals read in the said cluster. Then, the synthesis is written based on the views identified in the sample of proposals read. This protocol could be helpful to improve the scalability of participatory platforms with regard to the synthesis process.

However, we note that the interviewees also pointed out that our methodology to produce classifiers does not use extensively the categorization in sections of contributions proposed by their platforms. We acknowledge this remark since the categorizations proposed in both platforms studied contain valuable information concerning the content of proposals in each category. However, we consider that further use of the categorization in the training of our models would make it vulnerable to the risk that authors of proposals would classify them in an inappropriate category [91] such as the one placed at the top of the user interface [1], therefore reducing the reliability of our classifiers. This claim is supported by the purity metric computed for a section of RepNum based on its subcategories in Table 5.8. Nevertheless, categories defined in a platform can be used to narrow the scope of proposals processed for a specific use case, as proposed in the implementation of our first use case (Section 5.5.1).

Finally, we note that most of the discussion focused on the use case presented in

⁹A “suggestion box” participatory platform aims at gathering new ideas or vague projects from citizens. Certain ideas will then further developed and implemented by the organizing institution. These platforms typically implement very few features for citizens or institutions to give feedback to submitted ideas.

Section 5.5.1. We suppose that this focus stems from the fact that our interviewees are actively seeking to address the problem tackled by this first implementation, unlike the implementation proposed in the second use case. Indeed, the platform built by our interviewees currently does not implement NLP-based features for contributors. They consider that “the role of the user interface is to make it easier for citizens to find the right category for their proposals, and not necessarily to [help them] interact with others’ proposals” since many citizens engage with the platform only once¹⁰ to submit their position on a given debate. Based on this remark, we consider that our tool could be used during the writing process of a proposal to highlight proposals entailed with the text written, for example to reduce the number of duplicate proposals. Additionally, interviewees highlight the risk of manipulating participation with recommendation systems. While we agree that this risk should not be underestimated, we consider that our approach successfully addresses it by highlighting contradictory proposals and leveraging open-source models fine-tuned on publicly accessible and transparently listed corpora. This ensures that the fairness of the participatory initiative is not compromised by inadvertently under-representing specific idea communities.

5.7 Conclusion

This work presents research focused on the textual content of contributions in citizen participation platforms to help improve the experience of contributors and administrators on such platforms. We propose a methodology for classifying proposals as entailed, contradictory or neutral. We highlight the potential of such tool in two national citizen consultations. Many proposals are entailed or contradictory, hence there is a lot of potentially duplicated content and differences in opinion that should be addressed in priority. Moreover, the tool presented is shown to be capable of identifying idea communities in a citizen consultation and recommending contributions on the same topic as a given proposal while at the same time bringing new elements of discussion. We consider our work to be a significant step towards the creation of more viable large scale citizen participation platforms. We also address the problem of performing NLI on large texts by proposing alternative strategies explicitly designed for texts containing more than one sentence. Finally, there are many exciting future directions, for example, testing the NLI functionality during an ongoing citizen consultation or providing a taxonomy of contradiction types. Implementing automated extraction of characteristics in text contributions is a key element to improve online citizen collaboration, especially at large scale, both to reduce the information overload experienced by contributors and to help platform maintainers summarize the contributions.

Limitations. A potential limitation of our work is that our analysis is done on French text. It would be interesting to test the validity of our method for other languages. Cross-lingual studies help to identify whether our conclusions could be generalized to other languages. Unfortunately, we could not find a comparable source of citizen participation data regarding the dataset size. Consequently, the only method we identify to produce cross-lingual analysis on citizen participation

¹⁰25 % for the *RepNum* consultation according to this satisfaction survey [44].

contexts would be to translate the data sources we use in this study into other languages.

Chapter 6

General conclusions

Contents

6.1	Contributions	109
6.1.1	Analysis of a participatory platform’s design	109
6.1.2	Building a knowledge graph for citizen consultations	110
6.1.3	Identification of the different objectives and features implemented in participatory platforms	110
6.1.4	NLI for large-scale participation	110
6.2	Future research directions	111
6.2.1	Knowledge graphs	111
6.2.2	Qualitative analysis of interviews with stakeholders of citizen participation	111
6.2.3	Interaction of citizen participation platforms with other participation channels	111
6.2.4	Fine-grained identification of contradictions in proposals	112
6.2.5	Formalized workflow for online citizen participation	112

Given the rise in popularity of online citizen participation in the recent years, solving issues impacting participatory platforms has become crucial to improve inclusion of citizens in public decision making. In this thesis, we identify multiple issues faced by contemporary online participatory platforms. Based on the issues we identified and on previous works, we propose different solutions to improve online participation. These solutions are focused on providing better support for citizens collaboration and on handling citizen participation at large scale. In the following, we summarize the contributions produced during this thesis and propose potential directions for future research.

6.1 Contributions

6.1.1 Analysis of a participatory platform’s design

In Chapter 3, we identified biases in the design of the *République Numérique* participatory platform based on the interactions between citizens and contributions. We

first identified a centralization of replies around certain proposals. We found that this centralization is unrelated with the authors of these proposals and is more certainly due to the content of proposals themselves. Then, we describe the behavior of the most active users, identifying both a low degree of involvement of active voters in other types of contributions and the coexistence of both versatile and specialized active contributors. Finally, we highlight the biases in participation induced by the user interface design of the platform, and especially the positioning of the different proposals. These different issues identified in the *République Numérique* online consultation allow us to formulate five recommendations for improving participatory platforms.

6.1.2 Building a knowledge graph for citizen consultations

In Section 3.6 of this thesis, we propose a methodology for building a knowledge graph for online citizen participation. We implement this methodology using the contributions produced in the *République Numérique* platform. This knowledge graph-based representation can be used to allow interoperability between different participatory platforms and to enrich the contributions published by citizens using external resources. In the case of the *République Numérique* platform, we link proposals with different topics and with resources from the DBpedia knowledge base to improve understanding and facilitate the retrieval of proposals mentioning specific entities.

6.1.3 Identification of the different objectives and features implemented in participatory platforms

In Chapter 4, we propose a method to describe citizen participation platforms based on the objectives pursued by platforms maintainers and the features for interaction implemented. This description framework is based on a series of semi-structured interviews with different stakeholders of citizen participation in France and Spain. Based on the descriptions proposed and issues identified for the different categories, we formulate recommendations to improve citizen collaboration in participatory platforms and to reduce information overload faced by large-scale participatory platform users and administrators.

6.1.4 NLI for large-scale participation

In Chapter 5, we propose to facilitate online participation at large scale by identifying entailed and contradictory pairs of proposals in participatory platforms. We propose a methodology to produce classifiers for NLI tasks on citizen contributions. Then, we assess our methodology by implementing it on two different participatory platforms in French language. Finally, the classifiers produced are used to evaluate the relevance of our methodology for two different use cases: identification of clusters of entailed proposals and recommendation of proposals bringing new views or elements of contradiction to the readers of an initial proposal. The first use case is assessed using graph-related metrics. The second use case is assessed via an empirical evaluation conducted with independent testers.

6.2 Future research directions

6.2.1 Knowledge graphs

Online public consultations enable citizens to create proposals and participate in public debates and decision-making. However, given the large scale of these consultations – which cover various topics are presented in text format – managing and extracting meaningful insights from them can be a daunting task for both participatory platform administrators and institutions. Citizens may also find it challenging to locate information relevant to the consultation, identify proposals in a specific topic, and gain an overview of the relationship between existing proposals in a given consultation. To address these issues, we proposed a methodology in Section 3.6 for automatically constructing a consultation knowledge graph that can tackle these challenges. In the future, we plan to augment the consultation knowledge graph by incorporating additional data sources not stored as knowledge bases, such as the French National Assembly archive (<https://data.assemblee-nationale.fr>) to ease the creation of more informed and comprehensive proposals. Moreover, our topic modeling methodology could be improved by using automated labeling of topics [89], and empirically assessing the ability of DBpedia Spotlight to identify the most important words in proposals. We plan to perform this evaluation with a quantitative assessment of the relevance of the words annotated and of their associated annotation performed by independent testers.

6.2.2 Qualitative analysis of interviews with stakeholders of citizen participation

While we proposed a study of the different objectives pursued by platforms maintainers and the technical features they implement to achieve these objectives in Chapter 4, this analysis does not provide insights on the multiple objectives pursued and practices performed by citizens in participatory platforms. Previous studies have already tackled the topic of citizens' motivators to get involved in specific participatory initiatives [74, 116, 146]. However, we did not identify any studies focused on multiple and different participatory platforms. Consequently, future research could focus on conducting a series of semi-directed interviews with citizens involved in different citizen participation platforms to understand their motivations as well as their practices (e.g., do they interact with other citizens' contributions, do they produce proposals in collaboration with other citizens). These interviews will be the basis of a qualitative analysis focused on the differences and similarities identified between different citizen profiles.

6.2.3 Interaction of citizen participation platforms with other participation channels

Online citizen participation can take multiple shapes. In many cases, participatory platform are not the only place where an institution aims at taking into account citizens' opinions on certain issues. The most frequent example is probably social media services. However, there is little means for interaction between participatory platforms and external participation channels apart from link sharing. Therefore,

future research could focus on both comparative analysis of different online participation contexts (e.g., participatory platforms and social media) and implementation of means for interaction between these different contexts using social media interoperability [9, 10].

6.2.4 Fine-grained identification of contradictions in proposals

In Chapter 5 of this thesis, pairs of proposals are classified using three different labels: entailed, neutral and contradictory. However, pairs of texts can be in contradiction for multiple reasons (e.g., numeric disagreement, support of opposing views, etc.) and therefore could be treated differently. Therefore, implementing a more nuanced identification of contradictions in NLI models for citizen participation could help stakeholders understand the content of participatory platforms more accurately. However, the datasets used in Chapter 5 would not be relevant as they do not allow citizens to characterize their comments further than the pro / con stances presented. Producing large NLI datasets with more fine-grained labeling of contradictory pairs (e.g., numeric or ideological contradiction) in large texts would be an important first step toward a better characterization of contradictions in NLI. In the specific case of online citizen participation, this would be feasible by giving the possibility to citizens to position their comments in a more nuanced scale than simply “for” or “against” the associated proposal in future consultations.

6.2.5 Formalized workflow for online citizen participation

The contributions presented in this thesis are focused on specific aspects of online citizen participation and propose solutions to specific issues faced by their stakeholders. However, we did not address the challenge of incorporating these different contributions into a large-scale citizen participation platform. Consequently, a step towards the design of more reliable participatory platforms would be to formalize a new workflow for large-scale online participation based on our contributions, aiming at improving means for citizens collaboration and reducing issues already identified by us and other researchers. This workflow could be exemplified by the implementation of the associated technical features on an open-source participatory platform such as Decidim [51] or Consul [47].

Bibliography

- [1] William Aboucaya, Rafael Angarita, and Valerie Issarny. “Detecting Obstacles to Collaboration in an Online Participatory Democracy Platform: A Use-case Driven Analysis”. In: *44th International Conference on Software Engineering (ICSE’22). Proceedings of the 2nd International Workshop on Equitable Data and Technology*. FairWare ’22. Pittsburgh, Pennsylvania: Association for Computing Machinery, May 2022, pp. 25–33. ISBN: 9781450392921. DOI: 10.1145/3524491.3527307. URL: <https://hal.inria.fr/hal-03654027>.
- [2] William Aboucaya, Sonia Guehis, and Rafael Angarita. “Building Online Public Consultation Knowledge Graphs”. In: *Text2KG 2023: International Workshop on Knowledge Graph Generation from Text, Co-located with the ESWC 2023*. Ed. by CEUR-WS. Hersonissos (Crete), Greece, May 2023, p. 12. URL: <https://inria.hal.science/hal-04076743>.
- [3] B. Thomas Adler, Luca de Alfaro, Santiago M. Mola-Velasco, Paolo Rosso, and Andrew G. West. “Wikipedia Vandalism Detection: Combining Natural Language, Metadata, and Reputation Features”. In: *Computational Linguistics and Intelligent Text Processing*. Ed. by Alexander Gelbukh. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 277–288. ISBN: 978-3-642-19437-5.
- [4] Leman Akoglu, Hanghang Tong, and Danai Koutra. “Graph based anomaly detection and description: a survey”. In: *Data Mining and Knowledge Discovery* 29.3 (May 2015), pp. 626–688. ISSN: 1573-756X. DOI: 10.1007/s10618-014-0365-y. URL: <https://doi.org/10.1007/s10618-014-0365-y>.
- [5] Lucie Alexis, Christine Chevret-Castellani, and Sarah Labelle. “Enjeux sociopolitiques d’un nouvel ordre de discours autour des dispositifs de participation en lien avec la loi pour une «République numérique»”. In: *Cahiers du CRICIS [Méta]morphoses numériques de la culture et des médias : quelques pistes de réflexion*. Special Issue (2017), pp. 56–67.
- [6] Davide Alocci, Julien Mariethoz, Oliver Horlacher, Jerven T. Bolleman, Matthew P. Campbell, and Frederique Lisacek. “Property Graph vs RDF Triple Store: A Comparison on Glycan Substructure Search”. en. In: *PLOS ONE* 10.12 (2015). Publisher: Public Library of Science, e0144578. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0144578. URL: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0144578> (visited on 01/30/2023).
- [7] Isaac Alpizar-Chacon and Sergey Sosnovsky. “What’s in an index: Extracting domain-specific knowledge graphs from textbooks”. In: *Proceedings of the ACM Web Conference 2022*. 2022, pp. 966–976.

- [8] Luc André, Stéphane Martin, Gérald Oster, and Claudia-Lavinia Ignat. “Supporting Adaptable Granularity of Changes for Massive-scale Collaborative Editing”. In: *CollaborateCom - 9th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing - 2013*. Austin, United States, Oct. 2013. DOI: 10.4108/icst.collaboratecom.2013.254123. URL: <https://hal.inria.fr/hal-00903813>.
- [9] Rafael Angarita, Nikolaos Georgantas, and Valérie Issarny. “Social Middleware for Civic Engagement”. In: *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. 2019, pp. 1777–1786. DOI: 10.1109/ICDCS.2019.00176.
- [10] Rafael Angarita, Bruno Lefèvre, Shohreh Ahvar, Ehsan Ahvar, Nikolaos Georgantas, and Valérie Issarny. “Universal Social Network Bus: Toward the Federation of Heterogeneous Online Social Network Services”. In: *ACM Trans. Internet Technol.* 19.3 (Oct. 2019). ISSN: 1533-5399. DOI: 10.1145/3323333. URL: <https://doi.org/10.1145/3323333>.
- [11] Renzo Angles. “The Property Graph Database Model.” In: *AMW*. 2018.
- [12] Pablo Aragón, Andreas Kaltenbrunner, Antonio Calleja-López, Andrés Pereira, Arnau Monterde, Xabier E. Barandiaran, and Vicenç Gómez. “Deliberative Platform Design: The Case Study of the Online Discussions in Decidim Barcelona”. In: *Social Informatics*. Ed. by Giovanni Luca Ciampaglia, Afra Mashhadi, and Taha Yasserli. Cham: Springer International Publishing, 2017, pp. 277–287. ISBN: 978-3-319-67256-4.
- [13] Miguel Arana-Catania, Felix-Anselm Van Lier, Rob Procter, Nataliya Tkachenko, Yulan He, Arkaitz Zubiaga, and Maria Liakata. “Citizen Participation and Machine Learning for a Better Democracy”. In: *Digit. Gov.: Res. Pract.* 2.3 (July 2021). ISSN: 2691-199X. DOI: 10.1145/3452118. URL: <https://doi.org/10.1145/3452118>.
- [14] Miguel Arana-Catania, Rob Procter, Yulan He, and Maria Liakata. “Evaluation of Abstractive Summarisation Models with Machine Translation in Deliberative Processes”. In: *Proceedings of the Third Workshop on New Frontiers in Summarization*. 2021, pp. 57–64.
- [15] Sherry R. Arnstein. “A Ladder Of Citizen Participation”. In: *Journal of the American Institute of Planners* 35.4 (1969), pp. 216–224. DOI: 10.1080/01944366908977225. eprint: <https://doi.org/10.1080/01944366908977225>. URL: <https://doi.org/10.1080/01944366908977225>.
- [16] Dionysis Athanasopoulos. “Digital Ecclesia: Towards an Online Direct-Democracy Framework”. In: *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Society*. ICSE-SEIS '18. Gothenburg, Sweden: Association for Computing Machinery, 2018, pp. 91–94. ISBN: 9781450356619. DOI: 10.1145/3183428.3183432. URL: <https://doi.org/10.1145/3183428.3183432>.
- [17] ATILF, INIST, and LIPN. *PERCEO : un Projet d'Etiqueteur Robuste pour l'Écrit et pour l'Oral*. ORTOLANG (Open Resources and TOols for LANGuage) –www.ortolang.fr. 2012. URL: <https://hdl.handle.net/11403/perceo/v1>.

- [18] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. “DBpedia: A Nucleus for a Web of Open Data”. In: *The Semantic Web*. Ed. by Karl Aberer, Key-Sun Choi, Natasha Noy, Dean Allemang, Kyung-Il Lee, Lyndon Nixon, Jennifer Golbeck, Peter Mika, Diana Maynard, Riichiro Mizoguchi, Guus Schreiber, and Philippe Cudré-Mauroux. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 722–735. ISBN: 978-3-540-76298-0.
- [19] Sergio Bachiller, Lara Quijano-Sánchez, and Iván Cantador. “A Flexible and Lightweight Interactive Data Mining Tool to Visualize and Analyze Digital Citizen Participation Content”. In: *Proceedings of the 36th Annual ACM Symposium on Applied Computing*. SAC '21. Virtual Event, Republic of Korea: Association for Computing Machinery, 2021, pp. 413–416. ISBN: 9781450381048. DOI: 10.1145/3412841.3442081. URL: <https://doi.org/10.1145/3412841.3442081>.
- [20] Romain Badouard. “Combining Inclusion with Impact on the Decision? The Commission’s Online Consultation on the European Citizens’ Initiative”. In: *Is Europe listening to us?: Successes and failures of EU citizen consultations*. Ed. by Raphael Kies and Patrizia Nanz. Routledge, 2016, pp. 221–242. ISBN: 978-1-3155-8985-5. DOI: 10.4324/9781315589855.
- [21] João Barroca, Abhishek Shivkumar, Beatriz Quintino Ferreira, Evgeny Sherkhonov, and João Faria. “Enriching a fashion knowledge graph from product textual descriptions”. In: *arXiv preprint arXiv:2206.01087* (2022).
- [22] H.S. Baum. “Citizen Participation”. In: *International Encyclopedia of the Social & Behavioral Sciences*. Ed. by Neil J. Smelser and Paul B. Baltes. Oxford: Pergamon, 2001, pp. 1840–1846. ISBN: 978-0-08-043076-8. DOI: <https://doi.org/10.1016/B0-08-043076-7/04439-9>. URL: <http://www.sciencedirect.com/science/article/pii/B0080430767044399>.
- [23] Ulrich Beck. *La Société du Risque: Sur La voie d’une autre modernité*. Aubier, 2001. ISBN: 2081218887.
- [24] Lamjed Ben Jabeur, Lynda Tamine, and Mohand Boughanem. “Active Microbloggers: Identifying Influencers, Leaders and Discussers in Microblogging Networks”. In: *String Processing and Information Retrieval*. Ed. by Lilliana Calderón-Benavides, Cristina González-Caro, Edgar Chávez, and Nivio Ziviani. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 111–117. ISBN: 978-3-642-34109-0. DOI: 10.1007/978-3-642-34109-0_12. URL: https://doi.org/10.1007/978-3-642-34109-0_12.
- [25] Martin Biéri, Régis Chatellier, Antoine Courmont, Déborah Zribi, Hajar El Aoufir, Estelle Hary, Émile Gabrié, Ahlam Ammi, Pauline Faget, and Marie Leroux. “Making technological black boxes transparent: open source and open data”. In: *Civic tech, data and demos*. Vol. 7. Innovation and Foresight Reports. Paris, France: CNIL, Dec. 2019, pp. 23–24. URL: https://linc.cnil.fr/sites/default/files/atoms/files/cnil_ip_report_06_civic_tech_data_demos.pdf.
- [26] David M. Blei, Andrew Y. Ng, Michael I. Jordan, and John Lafferty. “Latent dirichlet allocation”. In: *Journal of Machine Learning Research* 3. Jan (2003), pp. 993–1022.

- [27] Loïc Blondiaux. “La force d’une notion floue”. In: *Le Nouvel esprit de la démocratie: Actualité de la Démocratie participative*. La République des Idées. Seuil, 2008, pp. 13–35. ISBN: 2020966751.
- [28] Kirsten Boehner and Carl DiSalvo. “Data, Design and Civics: An Exploratory Study of Civic Tech”. In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI ’16. San Jose, California, USA: Association for Computing Machinery, 2016, pp. 2970–2981. ISBN: 9781450333627.
- [29] Thomas Bonald, Nathan de Lara, Quentin Lutz, and Bertrand Charpentier. “Scikit-network: Graph Analysis in Python”. In: *Journal of Machine Learning Research* 21.185 (2020), pp. 1–6. URL: <http://jmlr.org/papers/v21/20-412.html>.
- [30] Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. “A large annotated corpus for learning natural language inference”. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, Sept. 2015, pp. 632–642. DOI: 10.18653/v1/D15-1075. URL: <https://aclanthology.org/D15-1075>.
- [31] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. “Language Models are Few-Shot Learners”. In: (2020). arXiv: 2005.14165 [cs.CL].
- [32] Axel Bruns and Jason Wilson. “Citizen Consultation from Above and Below : The Australian Perspective”. In: *Politics, Democracy and E-Government: Participation and Service Delivery*. Hershey, PA, USA: IGI Global, 2010, pp. 332–346. ISBN: 9781615209330. DOI: 10.4018/978-1-61520-933-0.ch020. URL: <http://services.igi-global.com/resolvedoi/resolve.aspx?doi=10.4018/978-1-61520-933-0.ch020>.
- [33] *Cahier des charges de la plateforme citoyenne Participer Strasbourg*. 2019. URL: <https://participer.strasbourg.eu/documents/2531722/3042639/Plan+PNU.pdf/740f2c9b-2fdf-df47-f252-f682af6f3397>.
- [34] Iván Cantador, María E. Cortés-Cediel, and Miriam Fernández. “Exploiting Open Data to analyze discussion and controversy in online citizen participation”. In: *Information Processing & Management* 57.5 (2020), p. 102301. ISSN: 0306-4573. DOI: <https://doi.org/10.1016/j.ipm.2020.102301>. URL: <https://www.sciencedirect.com/science/article/pii/S0306457320307962>.
- [35] Nico Carpentier. “Beyond the Ladder of Participation: An Analytical Toolkit for the Critical Analysis of Participatory Media Processes”. In: *Javnost - The Public* 23.1 (2016), pp. 70–88. DOI: 10.1080/13183222.2016.1149760. eprint: <https://doi.org/10.1080/13183222.2016.1149760>. URL: <https://doi.org/10.1080/13183222.2016.1149760>.

- [36] Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Spezia. “SemEval-2017 Task 1: Semantic Textual Similarity Multilingual and Crosslingual Focused Evaluation”. In: *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. Vancouver, Canada: Association for Computational Linguistics, Aug. 2017, pp. 1–14. DOI: 10.18653/v1/S17-2001. URL: <https://aclanthology.org/S17-2001>.
- [37] Scott Chacon and Ben Straub. *Pro Git*. Apress, 2014. URL: <https://github.com/progit/progit2/releases/download/2.1.194/progit.pdf>.
- [38] Kaiping Chen and Tanja Aitamurto. “Barriers for Crowd’s Impact in Crowdsourced Policymaking: Civic Data Overload and Filter Hierarchy”. In: *International Public Management Journal* 22.1 (2019), pp. 99–126. DOI: 10.1080/10967494.2018.1488780. eprint: <https://doi.org/10.1080/10967494.2018.1488780>. URL: <https://doi.org/10.1080/10967494.2018.1488780>.
- [39] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. *PaLM: Scaling Language Modeling with Pathways*. 2022. arXiv: 2204.02311 [cs.CL].
- [40] Mabi Clément. “Les conceptions de la participation citoyenne inscrites au sein d’un dispositif de concertation”. In: *Journées doctorales GIS Participation* (2011).
- [41] Jacob Cohen. “A Coefficient of Agreement for Nominal Scales”. In: *Educational and Psychological Measurement* 20.1 (1960), pp. 37–46. DOI: 10.1177/001316446002000104. eprint: <https://doi.org/10.1177/001316446002000104>. URL: <https://doi.org/10.1177/001316446002000104>.
- [42] Cap Collectif. *Bilan de la consultation nationale sur le projet de loi pour une république numérique*. Cap Collectif. Dec. 2015. URL: <https://www.republique-numerique.fr/media/default/0001/02/2377a996950cadd5a093dc9e3a621c0e1f9a19f5.pdf>.
- [43] Cap Collectif. *Consultation sur le projet de loi pour une République Numérique*. 2015. URL: <https://www.republique-numerique.fr/>.
- [44] Cap Collectif. *Résultats du questionnaire de satisfaction de la consultation nationale sur le projet de loi pour une république numérique*. Cap Collectif. Dec. 2015. URL: <https://www.republique-numerique.fr/media/default/0001/02/52587d44cde47eacdba246c5625232454c06e33a.pdf>.

- [45] Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. “XNLI: Evaluating Cross-lingual Sentence Representations”. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 2475–2485. DOI: 10.18653/v1/D18-1269. URL: <https://aclanthology.org/D18-1269>.
- [46] Michael Conover, Jacob Ratkiewicz, Matthew Francisco, Bruno Goncalves, Filippo Menczer, and Alessandro Flammini. “Political Polarization on Twitter”. In: *Proceedings of the International AAAI Conference on Web and Social Media* 5.1 (July 2011), pp. 89–96. URL: <https://ojs.aaai.org/index.php/ICWSM/article/view/14126>.
- [47] CONSUL. *The CONSUL Project*. URL: <https://consulproject.org/>.
- [48] Robert A. Dahl. *On Democracy*. Yale University Press, 1998. ISBN: 9780300254051. URL: <http://www.jstor.org/stable/j.ctv18zhcs4> (visited on 03/06/2023).
- [49] Joachim Daiber, Max Jakob, Chris Hokamp, and Pablo N. Mendes. “Improving Efficiency and Accuracy in Multilingual Entity Extraction”. In: *Proceedings of the 9th International Conference on Semantic Systems. I-SEMANTICS ’13*. Graz, Austria: Association for Computing Machinery, 2013, pp. 121–124. ISBN: 9781450319720. DOI: 10.1145/2506182.2506198. URL: <https://doi.org/10.1145/2506182.2506198>.
- [50] Marie-Catherine De Marneffe, Anna N. Rafferty, and Christopher D. Manning. “Finding contradictions in text”. In: *Proceedings of ACL-08: HLT*. Columbus, Ohio, USA: Association for Computational Linguistics, 2008, pp. 1039–1047.
- [51] Decidim. *Free open-source participatory democracy for cities and organizations*. URL: <https://decidim.org/>.
- [52] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. “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: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota: Association for Computational Linguistics, June 2019, pp. 4171–4186. DOI: 10.18653/v1/N19-1423. URL: <https://aclanthology.org/N19-1423>.
- [53] Alex J.A. Donkers, Dujuan Yang, and Nico H.G. Baken. “Linked Data for Smart Homes: Comparing RDF and Labeled Property Graphs”. English. In: *LDAC 2020 Linked Data in Architecture and Construction*. Ed. by María Poveda-Villalón, Ana Roxin, Kris McGlenn, and Pieter Pauwels. CEUR Workshop Proceedings. 8th Linked Data in Architecture and Construction Workshop, LDAC 2020, LDAC 2020 ; Conference date: 17-06-2020 Through 19-06-2020. CEUR-WS.org, 2020, pp. 23–36. URL: <http://linkedbuildingdata.net/ldac2020/>.
- [54] C. A. Ellis and S. J. Gibbs. “Concurrency Control in Groupware Systems”. In: *Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data*. SIGMOD ’89. Portland, Oregon, USA: Association for Computing Machinery, 1989, pp. 399–407. ISBN: 0897913175. DOI: 10.1145/67544.66963. URL: <https://doi.org/10.1145/67544.66963>.

- [55] Douglas C Engelbart and William K English. “A research center for augmenting human intellect”. In: *Proceedings of the December 9-11, 1968, fall joint computer conference, part I*. 1968, pp. 395–410.
- [56] Roberto Falanga. “The national participatory budget in Portugal: Opportunities and challenges for scaling up citizen participation in policymaking”. In: *Hope for democracy: 30 years of participatory budgeting worldwide* (2018), pp. 447–466.
- [57] Cristina Flesher Fominaya. *Democracy Reloaded: Inside Spain’s Political Laboratory from 15-M to Podemos*. Oxford University Press, June 2020. ISBN: 9780190099961. DOI: 10.1093/oso/9780190099961.001.0001. URL: <https://doi.org/10.1093/oso/9780190099961.001.0001>.
- [58] Nancy Fraser. “Rethinking the Public Sphere: A Contribution to the Critique of Actually Existing Democracy”. In: *Social Text* 25/26 (1990), pp. 56–80. ISSN: 01642472, 15271951. URL: <http://www.jstor.org/stable/466240> (visited on 03/06/2023).
- [59] Cornelia Fraune. “Gender matters: Women, renewable energy, and citizen participation in Germany”. In: *Energy Research & Social Science* 7 (2015), pp. 55–65. ISSN: 2214-6296. DOI: <https://doi.org/10.1016/j.erss.2015.02.005>. URL: <https://www.sciencedirect.com/science/article/pii/S2214629615000250>.
- [60] Bairan Fu, Wenming Zhang, Guangneng Hu, Xinyu Dai, Shujian Huang, and Jiajun Chen. “Dual Side Deep Context-Aware Modulation for Social Recommendation”. In: *Proceedings of the Web Conference 2021*. WWW ’21. Ljubljana, Slovenia: Association for Computing Machinery, 2021, pp. 2524–2534. ISBN: 9781450383127. DOI: 10.1145/3442381.3449940. URL: <https://doi.org/10.1145/3442381.3449940>.
- [61] Li Gan, Bin Gu, Sirkka Jarvenpaa, and Geng Yang. “Habit formation in online communities”. In: *ICIs 2009 Proceedings* (2009), p. 189.
- [62] Marcel Gauchet. “Crise dans la démocratie”. FR. In: *La revue lacanienne* 2.2 (2008), pp. 59–72. DOI: 10.3917/lr1.082.0059. URL: <https://doi.org/10.3917/lr1.082.0059>.
- [63] Erik Glaas, Mattias Hjerpe, Martin Karlson, and Tina-Simone Neset. “Visualization for Citizen Participation: User Perceptions on a Mainstreamed Online Participatory Tool and Its Usefulness for Climate Change Planning”. In: *Sustainability* 12.2 (2020). ISSN: 2071-1050. DOI: 10.3390/su12020705. URL: <https://www.mdpi.com/2071-1050/12/2/705>.
- [64] Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. “Knapsack Voting for Participatory Budgeting”. In: *ACM Trans. Econ. Comput.* 7.2 (July 2019). ISSN: 2167-8375. DOI: 10.1145/3340230. URL: <https://doi.org/10.1145/3340230>.
- [65] Max Goldman, Greg Little, and Robert C Miller. “Real-time collaborative coding in a web IDE”. In: *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM. 2011, pp. 155–164.

- [66] Gérard Grunberg. “Les « gilets jaunes » et la crise de la démocratie représentative”. FR. In: *Le Débat* 204.2 (2019), pp. 95–103. DOI: 10.3917/deba.204.0095. URL: <https://doi.org/10.3917/deba.204.0095>.
- [67] Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. “Annotation Artifacts in Natural Language Inference Data”. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*. New Orleans, Louisiana: Association for Computational Linguistics, June 2018, pp. 107–112. DOI: 10.18653/v1/N18-2017. URL: <https://aclanthology.org/N18-2017>.
- [68] Sara Hachem, Vivien Mallet, Raphaël Ventura, Animesh Pathak, Valérie Issarny, Pierre-Guillaume Raverdy, and Rajiv Bhatia. “Monitoring Noise Pollution Using the Urban Civics Middleware”. In: *2015 IEEE First International Conference on Big Data Computing Service and Applications*. 2015, pp. 52–61. DOI: 10.1109/BigDataService.2015.16.
- [69] Lile Hattori and Michele Lanza. “Syde: a tool for collaborative software development”. In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering-Volume 2*. 2010, pp. 235–238.
- [70] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia D’amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. “Knowledge Graphs”. In: *ACM Comput. Surv.* 54.4 (July 2021). ISSN: 0360-0300. DOI: 10.1145/3447772. URL: <https://doi.org/10.1145/3447772>.
- [71] J. Holston, V. Issarny, and C. Parra. “Engineering Software Assemblies for Participatory Democracy: The Participatory Budgeting Use Case”. In: *2016 IEEE/ACM 38th International Conference on Software Engineering Companion (ICSE-C)*. Austin, TX, USA: Institute of Electrical and Electronics Engineers, 2016, pp. 573–582.
- [72] Lester J Holtzblatt, Laurie E Damianos, and Daniel Weiss. “Factors impeding Wiki use in the enterprise: a case study”. In: *CHI’10 Extended Abstracts on Human Factors in Computing Systems*. 2010, pp. 4661–4676.
- [73] Anna Huang. “Similarity measures for text document clustering”. In: *Proceedings of the sixth new zealand computer science research student conference (NZCSRSC2008), Christchurch, New Zealand*. Vol. 4. 2008, pp. 9–56.
- [74] Shih-Wen Huang, Minhyang (Mia) Suh, Benjamin Mako Hill, and Gary Hsieh. “How Activists Are Both Born and Made: An Analysis of Users on Change.Org”. In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI ’15. Seoul, Republic of Korea: Association for Computing Machinery, 2015, pp. 211–220. ISBN: 9781450331456. DOI: 10.1145/2702123.2702559. URL: <https://doi.org/10.1145/2702123.2702559>.

- [75] Mario Ianniello, Silvia Iacuzzi, Paolo Fedele, and Luca Brusati. “Obstacles and solutions on the ladder of citizen participation: a systematic review”. In: *Public Management Review* 21.1 (2019), pp. 21–46. DOI: 10.1080/14719037.2018.1438499. eprint: <https://doi.org/10.1080/14719037.2018.1438499>. URL: <https://doi.org/10.1080/14719037.2018.1438499>.
- [76] Abdessamad Imine, Michaël Rusinowitch, Gérald Oster, and Pascal Molli. “Formal design and verification of operational transformation algorithms for copies convergence”. In: *Theoretical Computer Science* 351.2 (2006). Algebraic Methodology and Software Technology, pp. 167–183. ISSN: 0304-3975. DOI: <https://doi.org/10.1016/j.tcs.2005.09.066>. URL: <http://www.sciencedirect.com/science/article/pii/S030439750500616X>.
- [77] Insee. *Age pyramid in France (French)*. Insee. Jan. 2020. URL: <https://www.insee.fr/fr/statistiques/2381472>.
- [78] Insee. *Socio-professional category by gender and age in France (French)*. Insee. Feb. 2020. URL: <https://www.insee.fr/fr/statistiques/2489546>.
- [79] Hamed Jelodar, Yongli Wang, Chi Yuan, Xia Feng, Xiahui Jiang, Yanchao Li, and Liang Zhao. “Latent Dirichlet allocation (LDA) and topic modeling: models, applications, a survey”. In: *Multimedia Tools and Applications* 78 (2019), pp. 15169–15211.
- [80] Robert Johansen. *Groupware: Computer support for business teams*. The Free Press, 1988.
- [81] Rafael Kallis, Andrea Di Sorbo, Gerardo Canfora, and Sebastiano Panichella. “Ticket Tagger: Machine Learning Driven Issue Classification”. In: *2019 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 2019, pp. 406–409. DOI: 10.1109/ICSME.2019.00070.
- [82] Maged N. Kamel Boulos, Bernd Resch, David N. Crowley, John G. Breslin, Gunho Sohn, Russ Burtner, William A. Pike, Eduardo Jezierski, and Kuo-Yu Slayer Chuang. “Crowdsourcing, citizen sensing and sensor web technologies for public and environmental health surveillance and crisis management: trends, OGC standards and application examples”. In: *International Journal of Health Geographics* 10.1 (Dec. 2011), p. 67. ISSN: 1476-072X. DOI: 10.1186/1476-072X-10-67. URL: <https://doi.org/10.1186/1476-072X-10-67>.
- [83] Martin Karlsson. “Understanding Divergent Patterns of Political Discussion in Online Forums—Evidence from the European Citizens’ Consultations”. In: *Journal of Information Technology & Politics* 9.1 (2012), pp. 64–81. DOI: 10.1080/19331681.2012.635965. eprint: <https://doi.org/10.1080/19331681.2012.635965>. URL: <https://doi.org/10.1080/19331681.2012.635965>.
- [84] Christopher Kelty, Aaron Panofsky, Morgan Currie, Roderic Crooks, Seth Erickson, Patricia Garcia, Michael Wartenbe, and Stacy Wood. “Seven dimensions of contemporary participation disentangled”. In: *Journal of the Association for Information Science and Technology* 66.3 (2015), pp. 474–488. DOI: <https://doi.org/10.1002/asi.23202>. eprint: <https://asistdl.onlinelibrary.wiley.com/doi/pdf/10.1002/asi.23202>. URL: <https://asistdl.onlinelibrary.wiley.com/doi/abs/10.1002/asi.23202>.

- [85] J. A. Khan, Y. Xie, L. Liu, and L. Wen. “Analysis of Requirements-Related Arguments in User Forums”. In: *2019 IEEE 27th International Requirements Engineering Conference (RE)*. 2019, pp. 63–74. DOI: 10.1109/RE.2019.00018.
- [86] Antti Knutas, Victoria Palacin, Giovanni Maccani, and Markus Helfert. “Software Engineering in Civic Tech a Case Study about Code for Ireland”. In: *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Society*. ICSE-SEIS '19. Montreal, Quebec, Canada: IEEE Press, 2019, pp. 41–50. DOI: 10.1109/ICSE-SEIS.2019.14. URL: <https://doi.org/10.1109/ICSE-SEIS.2019.14>.
- [87] Travis Kriplean, Jonathan Morgan, Deen Freelon, Alan Borning, and Lance Bennett. “Supporting Reflective Public Thought with Considerit”. In: *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*. CSCW '12. Seattle, Washington, USA: Association for Computing Machinery, 2012, pp. 265–274. ISBN: 9781450310864. DOI: 10.1145/2145204.2145249. URL: <https://doi.org/10.1145/2145204.2145249>.
- [88] J. Richard Landis and Gary G. Koch. “The Measurement of Observer Agreement for Categorical Data”. In: *Biometrics* 33.1 (1977), pp. 159–174. ISSN: 0006341X, 15410420. URL: <http://www.jstor.org/stable/2529310> (visited on 06/22/2022).
- [89] Jey Han Lau, Karl Grieser, David Newman, and Timothy Baldwin. “Automatic Labelling of Topic Models”. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Portland, Oregon, USA: Association for Computational Linguistics, June 2011, pp. 1536–1545. URL: <https://aclanthology.org/P11-1154>.
- [90] Victoria Laurent. “« Les citoyens auteurs de la loi » : la consultation en ligne République numérique, une étude des médiations à l’œuvre dans le processus de construction de la loi”. 2021PA020076. Thèse de doctorat en sciences de l’information et de la communication. Paris 2, France, 2021. URL: <http://www.theses.fr/2021PA020076>.
- [91] Victoria Laurent, Romain Badouard, Clément Mabi, Cécile Méadel, and Guillaume Sire. “Les outils numériques gouvernementaux au service de la loi : Le design de la plateforme République numérique”. fr. In: *Terminal. Technologie de l’information, culture & société* 122 (June 2018). ISSN: 0997-5551. DOI: 10.4000/terminal.2071. URL: <http://journals.openedition.org/terminal/2071>.
- [92] John Lawrence and Chris Reed. “Argument Mining: A Survey”. In: *Computational Linguistics* 45.4 (Jan. 2020), pp. 765–818. ISSN: 0891-2017. DOI: 10.1162/coli_a_00364. eprint: https://direct.mit.edu/coli/article-pdf/45/4/765/1847520/coli_a_00364.pdf. URL: https://doi.org/10.1162/coli%5C_a%5C_00364.
- [93] Rémi Lefebvre. “Les Gilets jaunes et les exigences de la représentation politique”. In: *La Vie des idées* (Sept. 2019). ISSN: 2105-3030. URL: <https://laviedesidees.fr/Les-Gilets-jaunes-et-les-exigences-de-la-representation-politique.html>.

- [94] Bruno Lefevre and Valerie Issarny. “Matching technological & societal innovations: the social design of a mobile collaborative app for urban noise monitoring”. In: *2018 IEEE International Conference on Smart Computing (SMARTCOMP)*. IEEE. Taormina, Italy: Institute of Electrical and Electronics Engineers, 2018, pp. 33–40.
- [95] Daniel Lévy-Bruhl, Jean-Claude Desenclos, Sylvie Quelet, and François Bourdillon. “Extension of French vaccination mandates: from the recommendation of the Steering Committee of the Citizen Consultation on Vaccination to the law”. In: *Eurosurveillance* 23.17 (2018), pp. 18–00048.
- [96] Matteo Lissandrini, Davide Mottin, Katja Hose, and Torben Bach Pedersen. “Knowledge graph exploration systems: are we lost”. In: *Proc of the 12th Conference on Innovative Data Systems Research (CIDR)*. Vol. 22. 2022, pp. 10–13.
- [97] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. “Roberta: A robustly optimized bert pretraining approach”. In: *arXiv preprint arXiv:1907.11692* (2019).
- [98] G. Longford and C. Hurrell. “Online Citizen Consultation and Engagement in Canada”. In: *Encyclopedia of Digital Government*. Hershey, PA, USA: IGI Global, 2007, pp. 1261–1267. ISBN: 9781591407898. DOI: 10.4018/978-1-59140-789-8.ch191. URL: <http://services.igi-global.com/resolvedoi/resolve.aspx?doi=10.4018/978-1-59140-789-8.ch191>.
- [99] Martin Loughlin. “The Contemporary Crisis of Constitutional Democracy”. In: *Oxford Journal of Legal Studies* 39.2 (Mar. 2019), pp. 435–454. ISSN: 0143-6503. DOI: 10.1093/ojls/gqz005. eprint: <https://academic.oup.com/ojls/article-pdf/39/2/435/32374700/gqz005.pdf>. URL: <https://doi.org/10.1093/ojls/gqz005>.
- [100] Arthur Lupia. “Shortcuts Versus Encyclopedias: Information and Voting Behavior in California Insurance Reform Elections”. In: *American Political Science Review* 88.1 (1994), pp. 63–76. DOI: 10.2307/2944882.
- [101] Brendan Luyt. “Debating reliable sources: writing the history of the Vietnam War on Wikipedia”. In: *Journal of Documentation* 71.3 (Jan. 2015), pp. 440–455. ISSN: 0022-0418. DOI: 10.1108/JD-11-2013-0147. URL: <https://doi.org/10.1108/JD-11-2013-0147>.
- [102] Christopher D. Manning, Prabhakar Raghavan, and Hinric Schütze. “Evaluation of clustering”. In: *Introduction to information retrieval*. Vol. 16. 1. Cambridge University Press, 2009, pp. 356–360.
- [103] Frank Manola, Eric Miller, and Brian McBride. “RDF primer”. In: *W3C recommendation* 10.1-107 (2004), p. 6.
- [104] Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. “A SICK cure for the evaluation of compositional distributional semantic models”. In: *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*. Reykjavik, Iceland: European Language Resources Association (ELRA), May

- 2014, pp. 216–223. URL: http://www.lrec-conf.org/proceedings/lrec2014/pdf/363_Paper.pdf.
- [105] Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamé Seddah, and Benoît Sagot. “CamemBERT: a Tasty French Language Model”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Online: Association for Computational Linguistics, July 2020, pp. 7203–7219. DOI: 10.18653/v1/2020.acl-main.645. URL: <https://aclanthology.org/2020.acl-main.645>.
- [106] Tom McCoy, Ellie Pavlick, and Tal Linzen. “Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, July 2019, pp. 3428–3448. DOI: 10.18653/v1/P19-1334. URL: <https://aclanthology.org/P19-1334>.
- [107] Heleen L.P. Mees, Caroline J. Uittenbroek, Dries L.T. Hegger, and Peter P.J. Driessen. “From citizen participation to government participation: An exploration of the roles of local governments in community initiatives for climate change adaptation in the Netherlands”. In: *Environmental Policy and Governance* 29.3 (2019), pp. 198–208. DOI: <https://doi.org/10.1002/eet.1847>. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/eet.1847>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/eet.1847>.
- [108] Pablo N. Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. “DBpedia Spotlight: Shedding Light on the Web of Documents”. In: *Proceedings of the 7th International Conference on Semantic Systems*. I-Semantics ’11. Graz, Austria: Association for Computing Machinery, 2011, pp. 1–8. ISBN: 9781450306218. DOI: 10.1145/2063518.2063519. URL: <https://doi.org/10.1145/2063518.2063519>.
- [109] Pierre Mendès France. *La République Moderne*. Gallimard, 1962. ISBN: 2070350185.
- [110] Laura Michel and Loïc Blondiaux. “L’expertise en débat : jeux d’acteurs et conflits de savoirs autour d’un débat public local dans le Lot”. In: *Les constructions de l’action publique*. Ed. by Fabrizio Cantelli, Steve Jacob, Jean-Louis Genard, and Christian de Visscher. L’Harmattan, Jan. 2006, pp. 181–202.
- [111] Amita Misra, Brian Ecker, and Marilyn Walker. “Measuring the Similarity of Sentential Arguments in Dialogue”. In: *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. Los Angeles: Association for Computational Linguistics, Sept. 2016, pp. 276–287. DOI: 10.18653/v1/W16-3636. URL: <https://aclanthology.org/W16-3636>.
- [112] Chantal Mouffe. “Deliberative Democracy or Agonistic Pluralism?” In: *Social Research* 66.3 (1999), pp. 745–758. ISSN: 0037783X. URL: <http://www.jstor.org/stable/40971349> (visited on 03/06/2023).
- [113] David Nadeau and Satoshi Sekine. “A survey of named entity recognition and classification”. In: *Linguisticae Investigationes* 30.1 (2007), pp. 3–26.

- [114] William A Nagel. *Subversion Version Control: Using the Subversion Version Control System in Development Projects*. Vol. 14. Prentice Hall Professional Technical Reference, 2005.
- [115] Laura Naismith, B-H Lee, and Rachel M Pilkington. “Collaborative learning with a wiki: Differences in perceived usefulness in two contexts of use”. In: *Journal of Computer Assisted Learning* 27.3 (2011), pp. 228–242.
- [116] Mijail Naranjo-Zolotov, Tiago Oliveira, Frederico Cruz-Jesus, José Martins, Ramiro Gonçalves, Frederico Branco, and Nuno Xavier. “Examining social capital and individual motivators to explain the adoption of online citizen participation”. In: *Future Generation Computer Systems* 92 (2019), pp. 302–311. ISSN: 0167-739X. DOI: <https://doi.org/10.1016/j.future.2018.09.044>. URL: <https://www.sciencedirect.com/science/article/pii/S0167739X18313396>.
- [117] Matti Nelimarkka. “A Review of Research on Participation in Democratic Decision-Making Presented at SIGCHI Conferences. Toward an Improved Trading Zone Between Political Science and HCI”. In: *Proc. ACM Hum.-Comput. Interact.* 3.CSCW (Nov. 2019). DOI: 10.1145/3359241. URL: <https://doi.org/10.1145/3359241>.
- [118] Matthieu Nicolas, Victorien Elvinger, Gérald Oster, Claudia-Lavinia Ignat, and François Charoy. “MUTE: A Peer-to-Peer Web-based Real-time Collaborative Editor”. In: 2017.
- [119] Michele Notari. “How to use a Wiki in education: ‘Wiki based effective constructive learning’”. In: *Proceedings of the 2006 international symposium on Wikis*. 2006, pp. 131–132.
- [120] Natasha Noy, Yuqing Gao, Anshu Jain, Anant Narayanan, Alan Patterson, and Jamie Taylor. “Industry-scale Knowledge Graphs: Lessons and Challenges: Five diverse technology companies show how it’s done”. In: *Queue* 17.2 (2019), pp. 48–75.
- [121] *Octoverse 2022: The global developer community*. 2022. URL: <https://octoverse.github.com/2022/developer-community> (visited on 01/19/2023).
- [122] Claire Oger. “Fondements et formes de l’autorité en discours”. Habilitation à diriger des recherches. Paris, France: CELSA, Université Paris-Sorbonne, Dec. 2013.
- [123] Judith S. Olson, Dakuo Wang, Gary M. Olson, and Jingwen Zhang. “How People Write Together Now: Beginning the Investigation with Advanced Undergraduates in a Project Course”. In: *ACM Trans. Comput.-Hum. Interact.* 24.1 (Mar. 2017). ISSN: 1073-0516. DOI: 10.1145/3038919. URL: <https://doi.org/10.1145/3038919>.
- [124] Diler Öner. “Pre-service teachers’ experiences with Wiki: challenges of asynchronous collaboration”. In: *Proceedings of the 5th international Symposium on Wikis and Open Collaboration*. 2009, pp. 1–2.
- [125] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. *The PageRank Citation Ranking: Bringing Order to the Web*. Technical Report 1999-66. Stanford InfoLab, Nov. 1999. URL: <http://ilpubs.stanford.edu:8090/422/>.

- [126] Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. “A Decomposable Attention Model for Natural Language Inference”. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics, Nov. 2016, pp. 2249–2255. DOI: 10.18653/v1/D16-1244. URL: <https://aclanthology.org/D16-1244>.
- [127] Kevin Parker and Joseph Chao. “Wiki as a teaching tool”. In: *Interdisciplinary Journal of e-learning and Learning Objects* 3.1 (2007), pp. 57–72.
- [128] Cristhian Parra, Christelle Rohaut, Marianne Maeckelbergh, Valerie Issarny, and James Holston. “Expanding the Design Space of ICT for Participatory Budgeting”. In: *Proceedings of the 8th International Conference on Communities and Technologies*. C&T ’17. Troyes, France: Association for Computing Machinery, 2017, pp. 213–221. ISBN: 9781450348546. DOI: 10.1145/3083671.3083702. URL: <https://doi.org/10.1145/3083671.3083702>.
- [129] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. Vol. 32. Curran Associates, Inc., 2019. URL: <https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf>.
- [130] Carole Pateman. *Participation and Democratic Theory*. Cambridge University Press, 1970. DOI: 10.1017/CB09780511720444.
- [131] Ellie Pavlick, Matt Post, Ann Irvine, Dmitry Kachaev, and Chris Callison-Burch. “The Language Demographics of Amazon Mechanical Turk”. In: *Transactions of the Association for Computational Linguistics* 2 (Feb. 2014), pp. 79–92. ISSN: 2307-387X. DOI: 10.1162/tacl_a_00167. eprint: https://direct.mit.edu/tacl/article-pdf/doi/10.1162/tacl_a_00167/1566879/tacl_a_00167.pdf. URL: https://doi.org/10.1162/tacl%5C_a%5C_00167.
- [132] Oren Perez. “Complexity, Information Overload and Online Deliberation”. In: *Journal of Law and Policy for the Information Society* (2009), pp. 10–09.
- [133] Vicente Pina, Lourdes Torres, and Sonia Royo. “Comparing online with offline citizen engagement for climate change: Findings from Austria, Germany and Spain”. In: *Government Information Quarterly* 34.1 (2017). Open Innovation in the Public Sector, pp. 26–36. ISSN: 0740-624X. DOI: 10.1016/j.giq.2016.08.009. URL: <https://www.sciencedirect.com/science/article/pii/S0740624X16301496>.
- [134] The Computational Democracy Project. *Polis*. URL: <https://compdemocracy.org/Polis/>.
- [135] Robert D. Putnam. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster, 2000. ISBN: 978-0-7432-0304-3. URL: <http://bowlingalone.com/>.

- [136] Usha Nandini Raghavan, Réka Albert, and Soundar Kumara. “Near linear time algorithm to detect community structures in large-scale networks”. In: *Phys. Rev. E* 76 (3 Sept. 2007), p. 036106. DOI: 10.1103/PhysRevE.76.036106. URL: <https://link.aps.org/doi/10.1103/PhysRevE.76.036106>.
- [137] Nils Reimers and Iryna Gurevych. “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 3982–3992. DOI: 10.18653/v1/D19-1410. URL: <https://aclanthology.org/D19-1410>.
- [138] Nils Reimers and Iryna Gurevych. “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 3982–3992. DOI: 10.18653/v1/D19-1410. URL: <https://aclanthology.org/D19-1410>.
- [139] Nils Reimers and Iryna Gurevych. “Making Monolingual Sentence Embeddings Multilingual using Knowledge Distillation”. In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Nov. 2020. URL: <https://arxiv.org/abs/2004.09813>.
- [140] Ioannis Reklos and Albert Meroño-Peñuela. “MediCause: Causal Relation Modelling and Extraction from Medical Publications”. In: *CEUR Workshop Proceedings*. Vol. 3184. CEUR-WS. 2022, pp. 1–18.
- [141] Diego Rincon-Yanez and Sabrina Senatore. “FAIR Knowledge Graph construction from text, an approach applied to fictional novels”. In: (2022).
- [142] Ariella Rotramel, Rebecca Parmer, and Rose Oliveira. “Engaging Women’s History through Collaborative Archival Wikipedia Projects”. In: *The Journal of Interactive Technology and Pedagogy* 14 (Jan. 2019). ISSN: 2166-6245. URL: <https://jitp.commons.gc.cuny.edu/engaging-womens-history-through-collaborative-archival-wikipedia-projects>.
- [143] Jorge Saldivar, Florian Daniel, Luca Cernuzzi, and Fabio Casati. “Online Idea Management for Civic Engagement: A Study on the Benefits of Integration with Social Networking”. In: *Trans. Soc. Comput.* 2.1 (Jan. 2019). ISSN: 2469-7818. DOI: 10.1145/3284982. URL: <https://doi.org/10.1145/3284982>.
- [144] Víctor Sampedro and Josep Lobera. “The Spanish 15-M Movement: a consensual dissent?” In: *Journal of Spanish Cultural Studies* 15.1-2 (2014), pp. 61–80. DOI: 10.1080/14636204.2014.938466. eprint: <https://doi.org/10.1080/14636204.2014.938466>. URL: <https://doi.org/10.1080/14636204.2014.938466>.

- [145] Arthur D. Santana. “Virtuous or Vitriolic: The effect of anonymity on civility in online newspaper reader comment boards”. In: *Journalism Practice* 8.1 (2014), pp. 18–33. DOI: 10.1080/17512786.2013.813194. eprint: <https://doi.org/10.1080/17512786.2013.813194>. URL: <https://doi.org/10.1080/17512786.2013.813194>.
- [146] Lisa Schmidhuber, Dennis Hilgers, Thomas Gegenhuber, and Stefan Etzeltorfer. “The emergence of local open government: Determinants of citizen participation in online service reporting”. In: *Government Information Quarterly* 34.3 (2017), pp. 457–469. ISSN: 0740-624X. DOI: <https://doi.org/10.1016/j.giq.2017.07.001>. URL: <https://www.sciencedirect.com/science/article/pii/S0740624X16301502>.
- [147] Ronald Schroeter. “Engaging New Digital Locals with Interactive Urban Screens to Collaboratively Improve the City”. In: *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*. CSCW ’12. Seattle, Washington, USA: Association for Computing Machinery, 2012, pp. 227–236. ISBN: 9781450310864. DOI: 10.1145/2145204.2145239. URL: <https://doi.org/10.1145/2145204.2145239>.
- [148] Daniel Schulz and Jens Newig. “Assessing Online Consultation in Participatory Governance: Conceptual framework and a case study of a national sustainability-related consultation platform in Germany”. In: *Environmental Policy and Governance* 25.1 (2015), pp. 55–69. DOI: 10.1002/eet.1655. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/eet.1655>. URL: <https://doi.org/10.1002/eet.1655>.
- [149] Andrés Segura-Tinoco, Óscar G Borzdynski, and Iván Cantador. “An Argument-based Search Framework: Implementation on a Spanish Corpus in the E-Participation Domain”. In: *CIRCLE 2022: Proceedings of the 2nd Joint Conference of the Information Retrieval Communities in Europe*. Ed. by CEUR-WS. Samatan, Gers, France, 2022, p. 18.
- [150] Andrés Segura-Tinoco and Iván Cantador. “On the extraction and use of arguments in recommender systems: A case study in the e-participation domain”. In: *Proceedings of the 3rd edition of knowledge-aware and conversational recommender systems and the 5th edition of recommendation in complex environments workshops*. Ed. by CEUR-WS. Amsterdam, Netherlands, 2021.
- [151] Pnina Shachaf and Noriko Hara. “Beyond vandalism: Wikipedia trolls”. In: *Journal of Information Science* 36.3 (2010), pp. 357–370. DOI: 10.1177/0165551510365390. eprint: <https://doi.org/10.1177/0165551510365390>. URL: <https://doi.org/10.1177/0165551510365390>.
- [152] Sonali K Shah. “Motivation, governance, and the viability of hybrid forms in open source software development”. In: *Management science* 52.7 (2006), pp. 1000–1014.
- [153] Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. “Conflict-Free Replicated Data Types”. In: *Stabilization, Safety, and Security of Distributed Systems*. Ed. by Xavier Défago, Franck Petit, and Vincent Villain. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 386–400. ISBN: 978-3-642-24550-3.

- [154] Wei Shen, Jianyong Wang, and Jiawei Han. “Entity linking with a knowledge base: Issues, techniques, and solutions”. In: *IEEE Transactions on Knowledge and Data Engineering* 27.2 (2014), pp. 443–460.
- [155] Boaventura de Sousa Santos. “Participatory Budgeting in Porto Alegre: Toward a Redistributive Democracy”. In: *Politics & Society* 26.4 (1998), pp. 461–510. DOI: 10.1177/0032329298026004003. eprint: <https://doi.org/10.1177/0032329298026004003>. URL: <https://doi.org/10.1177/0032329298026004003>.
- [156] Karen Sparck Jones. “A Statistical interpretation of term specificity and its application in retrieval”. In: *Journal of Documentation* 28.1 (Jan. 1972), pp. 11–21. ISSN: 0022-0418. DOI: 10.1108/eb026526. URL: <https://doi.org/10.1108/eb026526>.
- [157] Manu Sporny, Dave Longley, Gregg Kellogg, Markus Lanthaler, and Niklas Lindström. “JSON-LD 1.1”. In: *W3C Recommendation, Jul* (2020).
- [158] Feng Su and Chris Beaumont. “Evaluating the use of a wiki for collaborative learning”. In: *Innovations in Education and Teaching International* 47.4 (2010), pp. 417–431.
- [159] R. Sumi, T. Yasserli, A. Rung, A. Kornai, and J. Kertesz. “Edit Wars in Wikipedia”. In: *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing*. 2011, pp. 724–727. DOI: 10.1109/PASSAT/SocialCom.2011.47.
- [160] D. Sun and C. Sun. “Context-Based Operational Transformation in Distributed Collaborative Editing Systems”. In: *IEEE Transactions on Parallel and Distributed Systems* 20.10 (2009), pp. 1454–1470.
- [161] Minna Tamper, Rafael Leal, Laura Sinikallio, Petri Leskinen, Jouni Tuominen, and Eero Hyvönen. “Extracting knowledge from parliamentary debates for studying political culture and language”. In: *Proceedings of the 1st International Workshop on Knowledge Graph Generation From Text and the 1st International Workshop on Modular Knowledge (TEXT2KG 2022 and MK2022)*. CEUR-WS. org. 2022.
- [162] Nathan TeBlunthuis, Aaron Shaw, and Benjamin Mako Hill. “Density Dependence Without Resource Partitioning: Population Ecology on Change.Org”. In: *Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. CSCW ’17 Companion. Portland, Oregon, USA: Association for Computing Machinery, 2017, pp. 323–326. ISBN: 9781450346887. DOI: 10.1145/3022198.3026358. URL: <https://doi.org/10.1145/3022198.3026358>.
- [163] Simon Tormey. “The Contemporary Crisis of Representative Democracy.” in: *Democratic Theory* 1.2 (2014), pp. 104–112. DOI: 10.3167/dt.2014.010211. URL: <https://www.berghahnjournals.com/view/journals/democratic-theory/1/2/dt010211.xml>.
- [164] Linus Torvalds and Junio Hamano. “Git: Fast version control system”. In: *URL http://git-scm.com* (2010).

- [165] Bogdan Vasilescu, Vladimir Filkov, and Alexander Serebrenik. “StackOverflow and GitHub: Associations between Software Development and Crowdsourced Knowledge”. In: *2013 International Conference on Social Computing*. 2013, pp. 188–195. DOI: 10.1109/SocialCom.2013.35.
- [166] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. “SciPy 1.0: fundamental algorithms for scientific computing in Python”. In: *Nature Methods* 17.3 (Mar. 2020), pp. 261–272. ISSN: 1548-7105. DOI: 10.1038/s41592-019-0686-2. URL: <https://doi.org/10.1038/s41592-019-0686-2>.
- [167] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. “SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems”. In: 32 (2019). Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett. URL: <https://proceedings.neurips.cc/paper/2019/file/4496bf24afe7fab6f046bf4923da8de6-Paper.pdf>.
- [168] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. “GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding”. In: *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*. Brussels, Belgium: Association for Computational Linguistics, Nov. 2018, pp. 353–355. DOI: 10.18653/v1/W18-5446. URL: <https://aclanthology.org/W18-5446>.
- [169] Chien-min Wang and David Turner. “Extending the wiki paradigm for use in the classroom”. In: *International Conference on Information Technology: Coding and Computing, 2004. Proceedings. ITCC 2004*. Vol. 1. IEEE. 2004, pp. 255–259.
- [170] Dakuo Wang, Haodan Tan, and Tun Lu. “Why Users Do Not Want to Write Together When They Are Writing Together: Users’ Rationales for Today’s Collaborative Writing Practices”. In: *Proc. ACM Hum.-Comput. Interact.* 1.CSCW (Dec. 2017). DOI: 10.1145/3134742. URL: <https://doi.org/10.1145/3134742>.
- [171] Jeremy K Ward, Florian Cafiero, Raphael Fretigny, James Colgrove, and Valérie Seror. “France’s citizen consultation on vaccination and the challenges of participatory democracy in health”. In: *Social Science & Medicine* 220 (2019), pp. 73–80.
- [172] S. Weiss, P. Urso, and P. Molli. “Logoot: A Scalable Optimistic Replication Algorithm for Collaborative Editing on P2P Networks”. In: *2009 29th IEEE International Conference on Distributed Computing Systems*. 2009, pp. 404–412.
- [173] S. Weiss, P. Urso, and P. Molli. “Logoot-Undo: Distributed Collaborative Editing System on P2P Networks”. In: *IEEE Transactions on Parallel and Distributed Systems* 21.8 (2010), pp. 1162–1174.

- [174] Jianshu Weng, Ee-Peng Lim, Jing Jiang, and Qi He. “TwitterRank: Finding Topic-Sensitive Influential Twitterers”. In: *Proceedings of the Third ACM International Conference on Web Search and Data Mining*. WSDM '10. New York, New York, USA: Association for Computing Machinery, 2010, pp. 261–270. ISBN: 9781605588896.
- [175] Manjula Wijewickrema, Vivien Petras, and Naomal Dias. “Selecting a text similarity measure for a content-based recommender system”. In: *The Electronic Library* 37.3 (Jan. 2019), pp. 506–527. ISSN: 0264-0473. DOI: 10.1108/EL-08-2018-0165. URL: <https://doi.org/10.1108/EL-08-2018-0165>.
- [176] Alexander Wilson and Mark Tewdwr-Jones. “Let’s draw and talk about urban change: Deploying digital technology to encourage citizen participation in urban planning”. In: *Environment and Planning B: Urban Analytics and City Science* 47.9 (2020), pp. 1588–1604. DOI: 10.4000/terminal.2071.
- [177] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. “Transformers: State-of-the-Art Natural Language Processing”. In: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Online: Association for Computational Linguistics, Oct. 2020, pp. 38–45. URL: <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- [178] Scott Wright. “‘Success’ and online political participation: The case of Downing Street E-petitions”. In: *Information, Communication & Society* 19.6 (2016), pp. 843–857. DOI: 10.1080/1369118X.2015.1080285. eprint: <https://doi.org/10.1080/1369118X.2015.1080285>. URL: <https://doi.org/10.1080/1369118X.2015.1080285>.
- [179] W. Xing and A. Ghorbani. “Weighted PageRank algorithm”. In: *Proceedings. Second Annual Conference on Communication Networks and Services Research, 2004*. Fredericton, NB, Canada: Institute of Electrical and Electronics Engineers, 2004, pp. 305–314.
- [180] Taha Yasseri, Robert Sumi, András Rung, András Kornai, and János Kertész. “Dynamics of conflicts in Wikipedia”. In: *PloS one* 7.6 (2012), e38869.
- [181] Bo Yu and Guoray Cai. “Facilitating Participatory Decision-Making in Local Communities through Map-Based Online Discussion”. In: *Proceedings of the Fourth International Conference on Communities and Technologies*. C&T '09. University Park, PA, USA: Association for Computing Machinery, 2009, pp. 215–224. ISBN: 9781605587134. DOI: 10.1145/1556460.1556492. URL: <https://doi.org/10.1145/1556460.1556492>.
- [182] Mijail Naranjo Zolotov, Tiago Oliveira, and Sven Casteleyn. “Continued Intention to Use Online Participatory Budgeting: The Effect of Empowerment and Habit”. In: *Proceedings of the 11th International Conference on Theory and Practice of Electronic Governance*. ICEGOV '18. Galway, Ireland: Association for Computing Machinery, 2018, pp. 209–216. ISBN: 9781450354219. DOI: 10.1145/3209415.3209461. URL: <https://doi.org/10.1145/3209415.3209461>.