Social gender and derivational morphology:

A distributional study of the gendered import of learned morphology in French

Marine Wauquier & Olivier Bonami

Université de Paris, LLF & CNRS

Context

- French suffixes -*euse* and -*rice* are clear morphological rivals
  - Feminine agent nouns
    - *danser* 'dance' > *danseuse* 'female dancer'
    - *rédiger* 'write' > *rédactrice* 'female author'
  - Instrument nouns
    - *agrafer* 'staple' > *agrafeuse* 'stapler'
    - *excaver* 'excavate' > *excavatrice* 'excavator'
- Literature has put forth semantic differences with respect to the derived agent nouns (Lenoble-Pinson, 2008; Dawes, 2003, among others)
  - Denotation of low-level professions for -*euse*
    - *serveuse* 'waitress', *entraîneuse* 'barmaid'
  - Denotation of more socially valued positions for -*rice*
    - *directrice* 'female manager', *sénatrice* 'female senator'
- Recent studies show the distributional relevance of this distinction (Wauquier et al., 2020)

Further differences

- Both suffixes are not strictly identical with respect to their morphological construction
  - -*rice* originates from learned vocabulary (Rainer and Buridant, 2015) similarly to -*ion* and -*if* among others
- This distinction is hardly ever taken into account
  - It's is also at stake for the formation of masculine agent nouns
    - Learned -*eur* attaches to the same stems as -*ion* ('hidden stem', see Bonami et al. 2009) – *fondateur* 'founder', *fondation* 'foundation' from *fonder* 'found'
    - Nonlearned -*eur* attaches to the same nonlearned stems as -*age* – *dresseur* 'trainer', *dressage* 'training' from *dresser* 'train'

Research question

- To what extent does the semantic distinction between -*euse* and -*rice* follow from their morphological specificity, i.e. their status as learned vs. nonlearned formation?
  - If it does, we expect the differences between learned and nonlearned formations to be parallel for masculine and feminine agent nouns
  - By extension, parallel effects should be found for action nouns in -*ion* and -*age*
- We combine distributional and computational approaches to provide an empirical assessment of these hypotheses.

(Non)learned formations in a cube

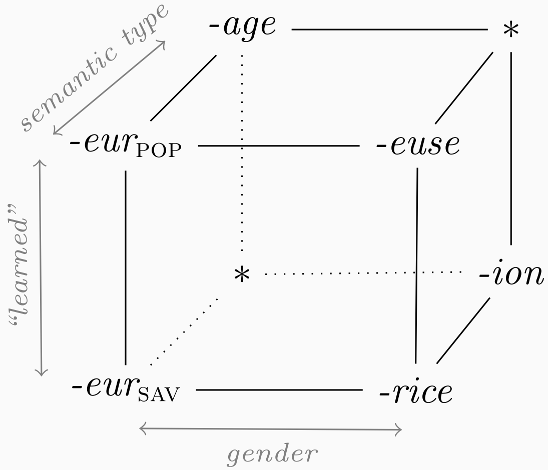

(Non)learned formations in a cube

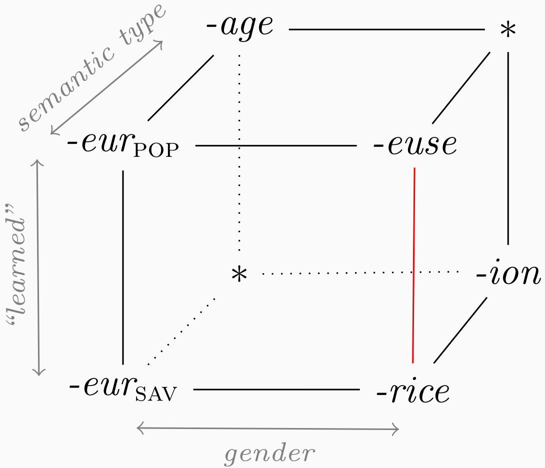

(Non)learned formations in a cube

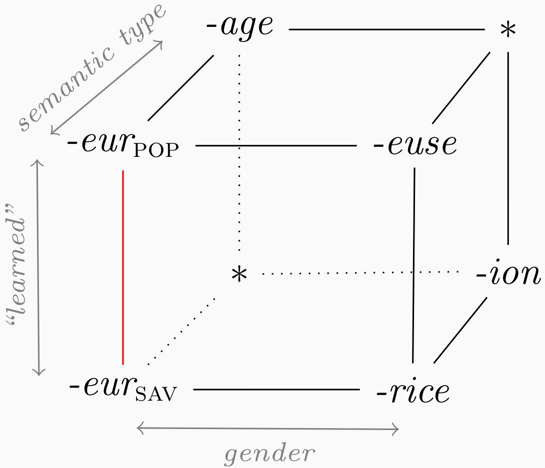

(Non)learned formations in a cube

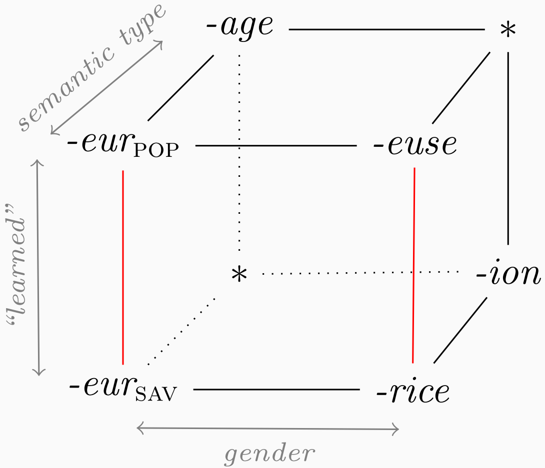

(Non)learned formations in a cube

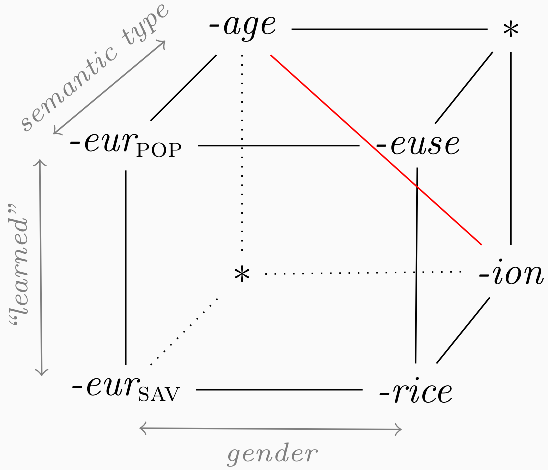

(Non)learned formations in a cube

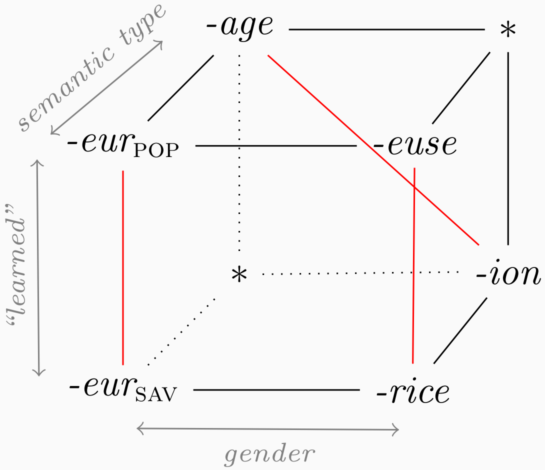

Distributional semantics in a nutshell

- Distributional semantics is grounded in the hypothesis that differences in word meanings are reflected by differences in distribution, i.e. the contexts in which they appear.
- The meaning of words in a given corpus is represented by vectors computed based on their co-occurrences
- The semantic similarity between two words is assessed by the distributional proximity of their vectors on a scale from 0 to 1.

Context \ Target	*dog*	*cat*	*car*
kibble	6	9	1
gas	2	1	8

Table 1: Co-occurences matrix of 3 targets

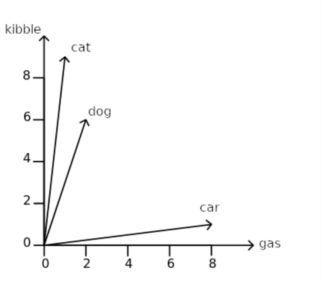

Figure 1: Targets' vectors

Our vectors

- Our vectors are trained
  - with Word2Vec (Mikolov et al., 2013) basic parameters
    - CBOW, 100 dimensions, Negative Sampling, frequency threshold of 5, window of 5
  - On the FrCoW corpus (Schäfer and Bildhauer, 2012; Schäfer, 2015) which was:
    - Lemmatized – e.g. *dînera* → `dîner_ver`
    - Tagged – e.g. *un dîner* → `un_art dîner_nom`
    - Carefully gender-neutralized – e.g. *du* → `de_prep le_art`

Intrinsic classification task

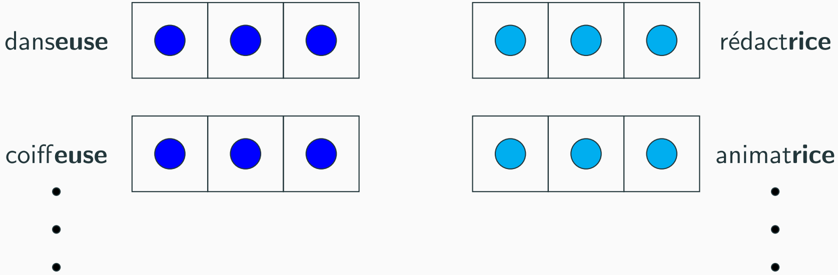

Intrinsic classification task

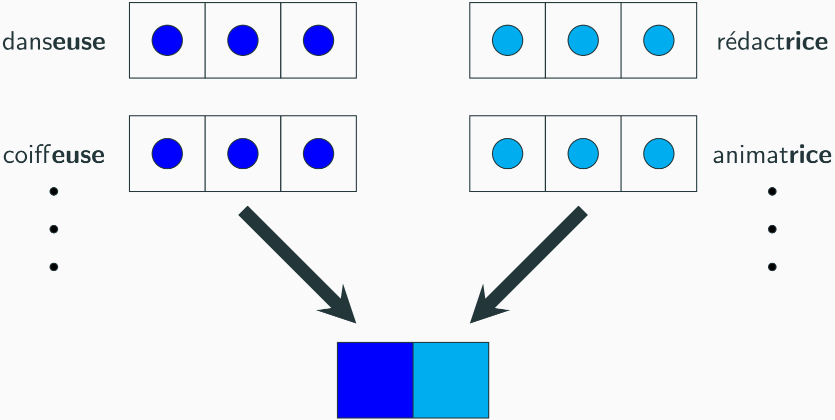

Intrinsic classification task

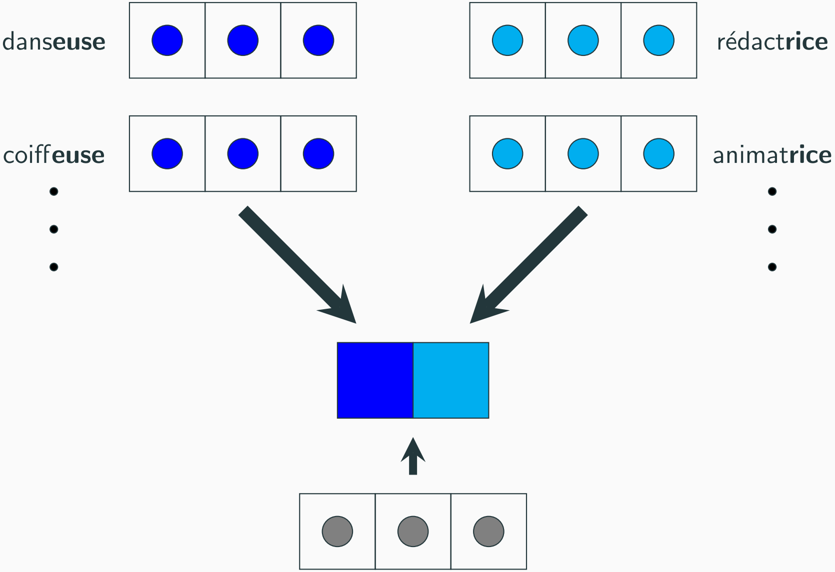

Intrinsic classification task

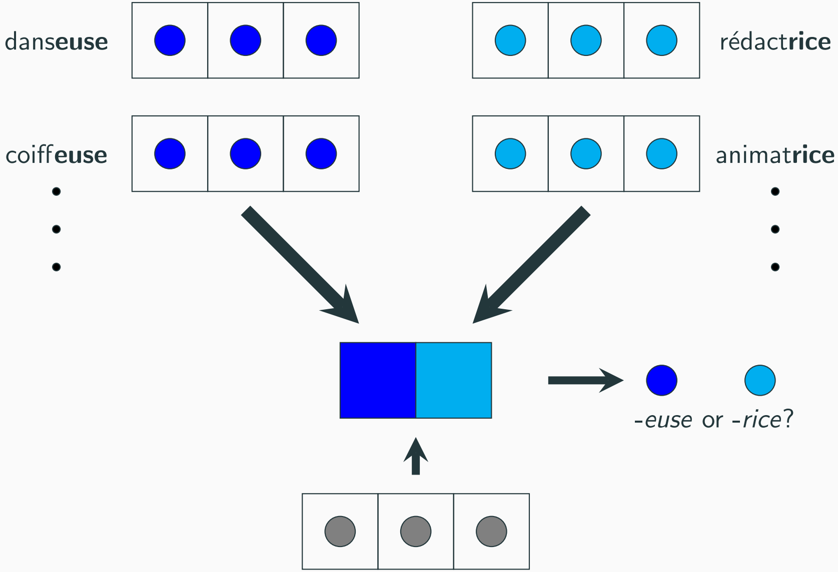

Extrinsic prediction task

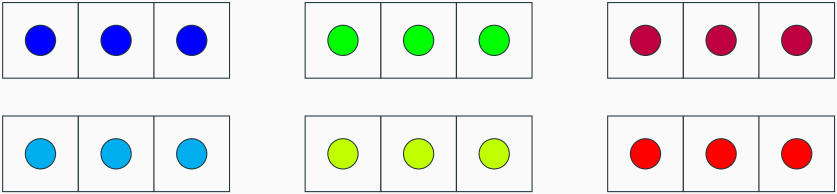

Extrinsic prediction task

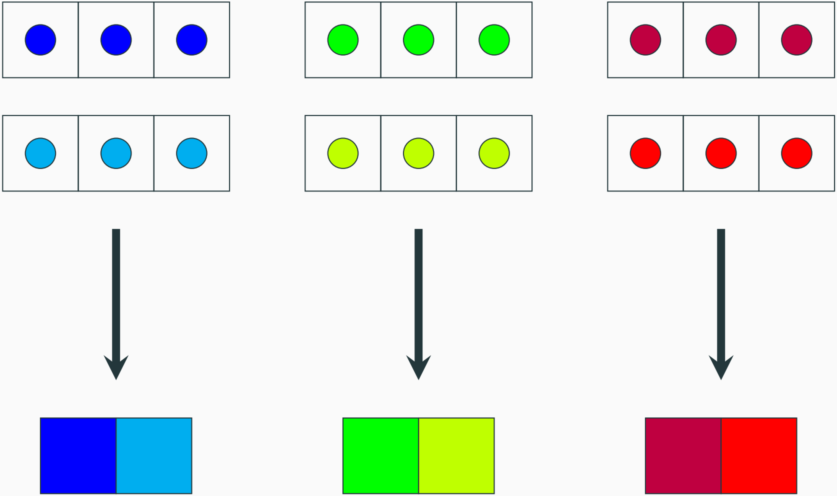

Extrinsic prediction task

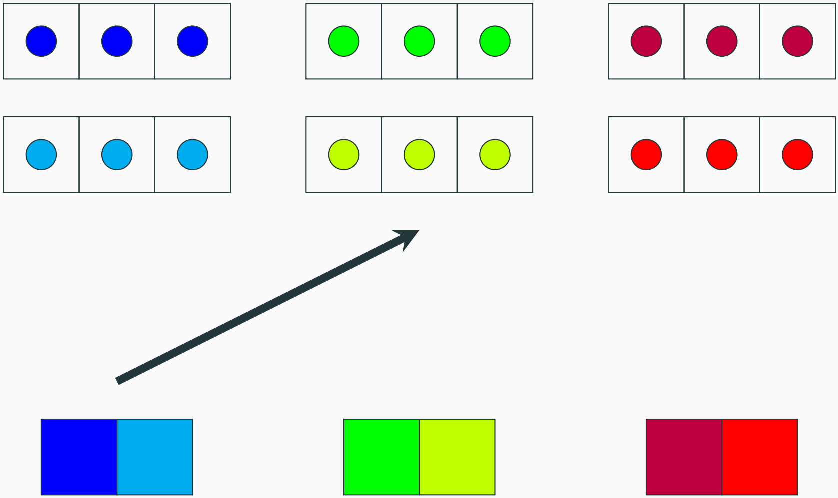

Extrinsic prediction task

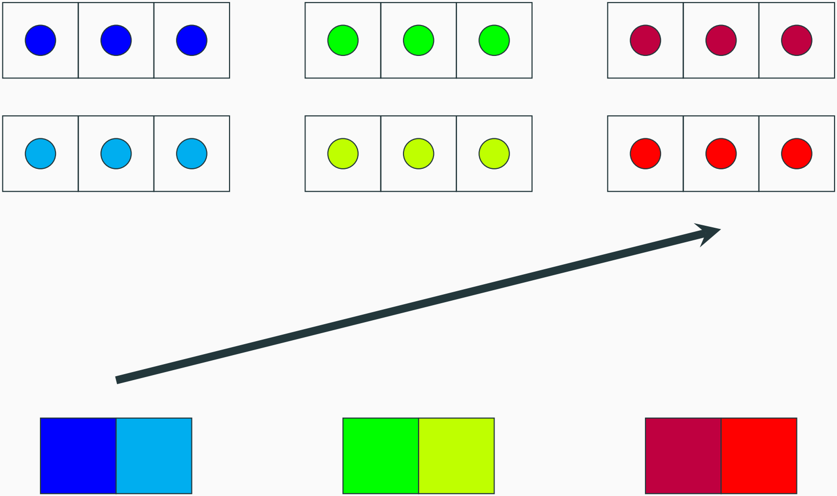

Extrinsic prediction task

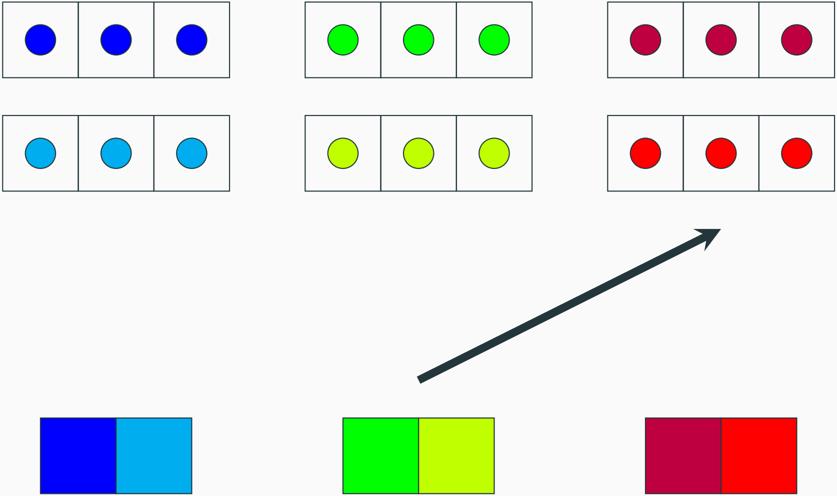

Extrinsic prediction task

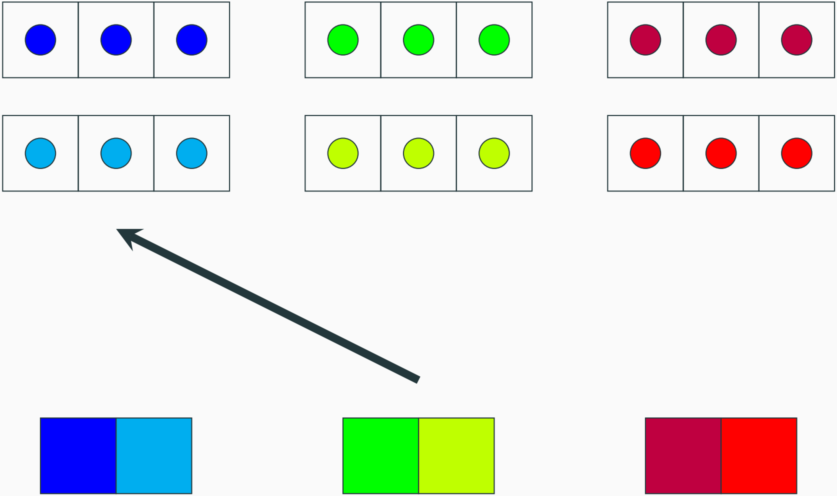

Extrinsic prediction task

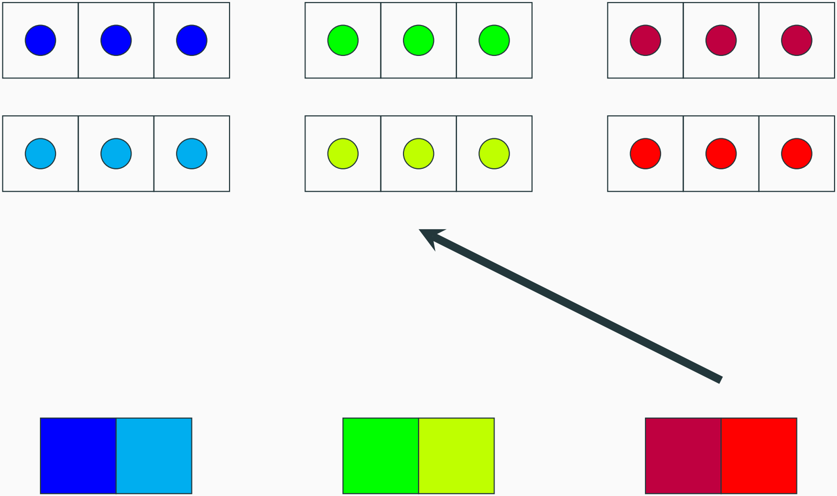

Extrinsic prediction task

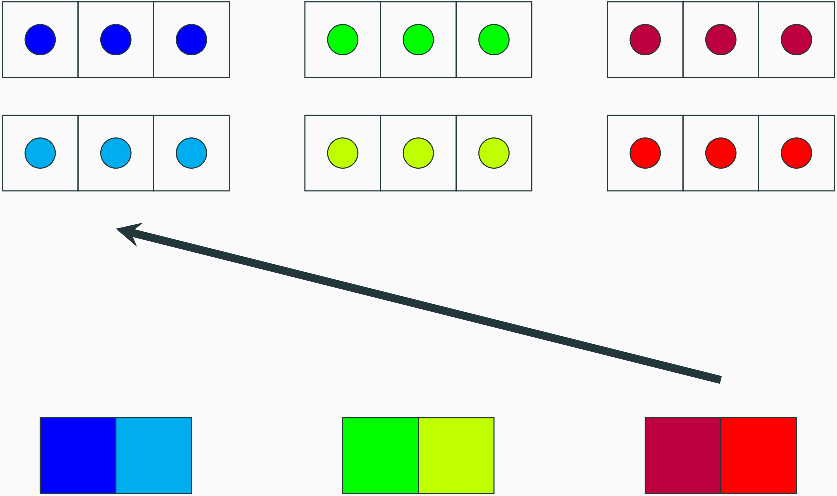

Datasets

- We build three dataset of feminine agent nouns (AGF), masculine agent nouns (AGM) and action nouns (ACT) with the (non)learned alternation
  - All agent nouns were filtered to exclude polysemy
  - Only nouns with a frequency of 50 or more in the FrCow corpus are kept

	Learned	Nonlearned
AGF (*rice* vs. *euse*)	158	301
AGM	141	462
ACT (*ion* vs. *age*)	750	629

Table 2: Description of our datasets

Methodology

- We use gradient boosting (Friedman, 2001; Mason et al., 2000) applied to decision trees as our binary classification method
  - 500 estimators, max depth of 2, deviance loss function
- The classifiers are trained on 282 nouns (141 of each suffix)
  - We randomly sub-sample our datasets based on the smallest one
- The performance of the classifiers are evaluated by means of
  - a 10-fold cross-validation for intrinsic predictions
  - a confusion matrix for extrinsic predictions, where suffixes are matched based on their (non)learned status
    - e.g. when the classifier trained on action nouns to predict -*ion* is applied to feminine agent nouns, -*rice* nouns labeled as -*ion* are considered as true positive

Expectations

- Because we hypothesize that the (non)learned feature is a distinctive but shared feature between all three datasets
  - We expect intrinsic predictions to get good results
  - We expect extrinsic predictions to perform better than the baseline, even similarly to intrinsic predictions

Intrinsic prediction results

- All three classifiers reach relatively high accuracy

	Accuracy	Confidence interval
AGF	80%	75.1 - 84.5
AGM	77%	72 - 81.9
ACT	83%	78.6 - 87.4

Table 3: Performance of intrinsic classifiers

- Despite the small training set, distribution is distinctive enough between learned and nonlearned nouns within each dataset
    - Yet it does not entail that all three datasets vary similarly

Extrinsic prediction results

- Intrinsic and extrinsic predictions lead to very similar accuracy

	Test data		
Training data	AGF	AGM	ACT
AGF	80%	77%	79%
AGM	77%	77%	82%
ACT	76%	79%	83%

Table 4: Accuracy of the three classifiers applied to the three datasets

- The learned vs. nonlearned distinction is equally predicted by all models on all datasets
- This distinction is stable and relevant in the distribution of all three categories of nouns

Qualitative assessment

- It remains to be seen whether this learned vs. nonlearned distinction is link to the observed difference in connotations between -*euse* and -*rice*
    - If so, a difference should be observed for the masculine agent nouns
- We rely on a qualitative assessment of their distribution
    - We build an average representation (*centroid*) for each suffix which we qualify by means of its 100 nearest neighbors

Centroid and neighbors

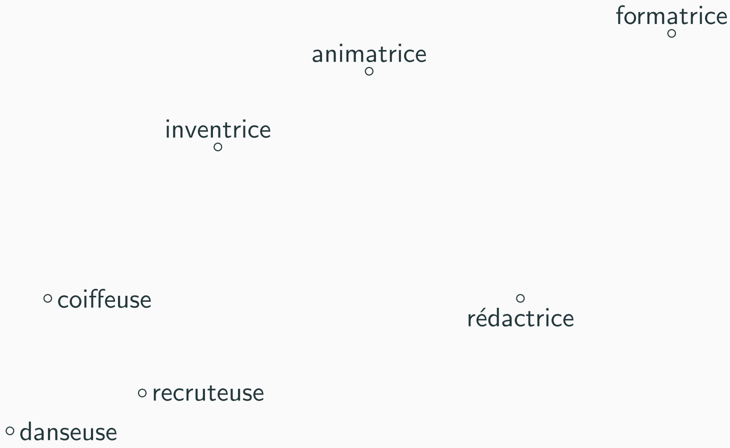

Centroid and neighbors

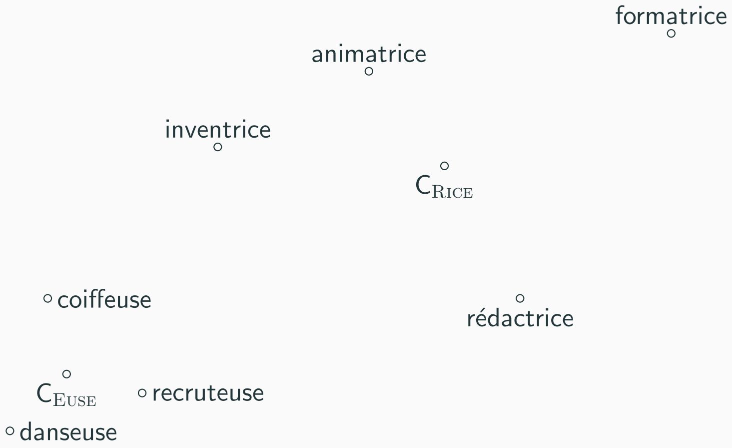

Centroid and neighbors

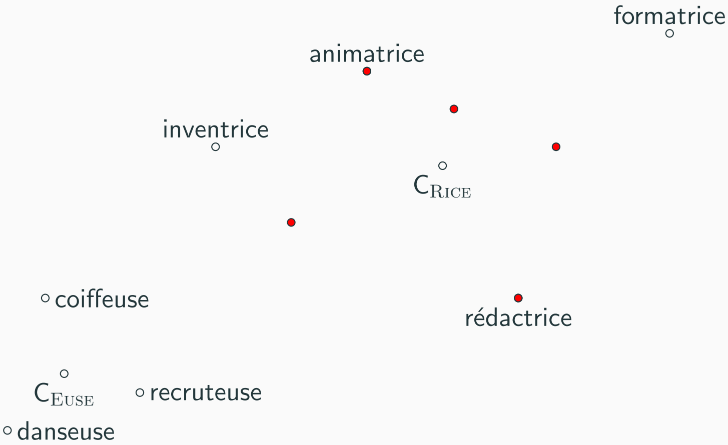

Semantic properties of the neighbors i

- Learned and nonlearned centroids have distinctive neighborhood
  - AGF centroids only share 8 neighbors out of 100
  - AGM centroids only share 3 neighbors out of 100

Semantic properties of the neighbors i

- Learned and nonlearned centroids have distinctive neighborhood
  - AGF centroids only share 8 neighbors out of 100
  - AGM centroids only share 3 neighbors out of 100
- Learned centroids neighborhoods contain a much lower proportion of negatively valued neighbors than nonlearned centroids neighborhoods
  - Positively valued neighbors
    - *dirigeante* 'female leader', *chirurgienne* 'female surgeon', *avocate* 'female lawyer', *poétesse* 'female poet'
    - *érudit* 'scholar', *académicien* 'academician', *orateur* 'orator', *intellectuel* 'intellectual'
  - Neutral neighbors
    - *camerawoman* 'camerawoman', *formatrice* 'female tutor', *animatrice* 'female presenter'
    - *exécutant* 'subordinate', *journaliste* 'journalist', *comptable* 'accountant', *contributeur* 'contributor'

Semantic properties of the neighbors ii

- The axiological properties of the nonlearned centroids neighbors vary with regard to gender
  - With respect to the feminine agent nouns (AGF), the neighbors involve
    - **Sexuality** – *nymphomane* 'nymphomaniac', *tapineuse* 'prostitute', *catin* 'harlot', *allumeuse* 'tease', *hardeuse* 'pornographic film actress'
    - **Physical characterization** – *laideron* 'plain Jane', *monstresse* 'monstress', *midinette* 'starry-eyed girl'
  - With respect to masculine agent nouns (AGM), they involve
    - **Sexuality** – *dragueur* 'womanizer', *séducteur* 'seducer'
    - **Behavioral characterization** – *tire-au-flanc* 'slacker', *poivrot* 'drunkard', *rustre* 'lout', *paresseux* 'idler'
    - **Criminal activities** – *truand* 'gangster', *voleur* 'thief', *malandrin* 'brigand'

Conclusion i

- We showed that the learned vs. nonlearned distinction instantiated with French feminine and masculine agent nouns and action nouns was distributionally identifiable
    - This feature is discriminative within each category
    - This feature is stable across all three categories
- We highlighted the semantic correlate of the learned vs. nonlearned distinction that varies depending on the morphosemantic type
    - It is instantiated in terms of axiological valence for agent nouns
        - The axiological values differ with respect to the gender and the (non)learned status
    - It implements a contrast between intellectual and technical domains of reference for action nouns (Wauquier et al., 2020)
    - Whether these differences comes particular sociolinguistic circumstances in which learned formations entered the language is yet to be confirmed

Conclusion ii

- We propose an innovative method to scrutinize the manifestation of a given feature within the lexicon
    - This methodology can be extended to other paradigmatic structures and features

Thanks

- Institutions
  - ANR Project *Demonext* (PI Fiammetta Namer)
  - Laboratoire de linguistique formelle (Université de Paris & CNRS)

References

Bonami, O., Boyé, G., and Kerleroux, F. (2009). L'allomorphie radicale et la relation flexion-construction. In Fradin, B., Kerleroux, F., and Plénat, M., editors, *Aperçus de morphologie du français*, pages 103–125. Presses de l'Université de Vincennes, Saint-Denis.

Dawes, E. (2003). La féminisation des titres et fonctions dans la francophonie: de la morphologie à l'idéologie. *Ethnologies*, 25(2):195–213.

Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.

Lenoble-Pinson, M. (2008). Mettre au féminin les noms de métier: résistances culturelles et sociolinguistiques. *Le français aujourd'hui*, (4):73–79.

Mason, L., Baxter, J., Bartlett, P. L., and Frean, M. R. (2000). Boosting algorithms as gradient descent. In *Advances in neural information processing systems*, pages 512–518.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781.

Rainer, F. and Buridant, C. (2015). From old french to modern french. In Muller, P., Ohnheiser, I., Olsen, S., and Rainer, F., editors, *Word-Formation: an international handbook of the languages of Europe*, volume 3, pages 1975–2000. De Gruyter Mouton, Berlin/Boston.

Schäfer, R. (2015). Processing and querying large web corpora with the COW14 architecture. In *Proceedings of Challenges in the Management of Large Corpora*, pages 28–34.

Schäfer, R. and Bildhauer, F. (2012). Building large corpora from the web using a new efficient tool chain. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation*, pages 486–493.

Wauquier, M., Hathout, N., and Fabre, C. (2020). Contributions of distributional semantics to the semantic study of French morphologically derived agent nouns. In Audring, J., Koutsoukos, N., and Manouilidou, C., editors, *Rules, patterns, schemas and analogy, MMM12 Online Proceedings*, volume 12, pages 111–121.

Inflectional cube