AI Clinical Opportunities: Opportunities and Pitfalls

Translating minimally-supervised NLP to interpretable clinical AI

John Kang, MD, PhD

@DrJohnKang

AAPM 2022
JULY 10–14 | WASHINGTON, DC
64TH ANNUAL MEETING & EXHIBITION
CELEBRATING MEDICAL PHYSICS
TRANSFORMING HUMAN HEALTH

UNIVERSITY of WASHINGTON

UW Medicine
DEPARTMENT OF RADIATION ONCOLOGY

Potential conflicts of interest

- Family member works for Change Healthcare

Outline: a tour of how NLP has leveraged unstructured data, unsupervised learning, and transfer learning to extract information from patient records

1. Supervised models for clear outcomes using structured data
2. Supervised models for clear outcomes using unstructured data
3. Supervised models for unclear outcomes using unstructured data
4. Supervised models for unclear outcomes using pre-trained unstructured data
5. Unsupervised models for unclear outcomes using pre-trained unstructured data

Real world trials using supervised learning on structured data

Research

JAMA Oncology | Original Investigation

Validation of a Machine Learning Algorithm to Predict 180-Day Mortality for Outpatients With Cancer

Christopher R. Manz, MD; Jinbo Chen, PhD; Manqing Liu, MHS; Corey Chivers, PhD; Susan Harkness Regli, PhD; Jennifer Braun, MHA; Michael Draugelis, MS; C. William Hanson, MD; Lawrence N. Shulman, MD; Lynn M. Schuchter, MD; Nina O'Connor, MD; Justin E. Bekelman, MD; Mitesh S. Patel, MD, MBA; Ravi B. Parikh, MD, MPP

Figure 1. Overall 180-Day Mortality of Patients Considered High vs Low Risk as Identified by a Machine Learning Algorithm

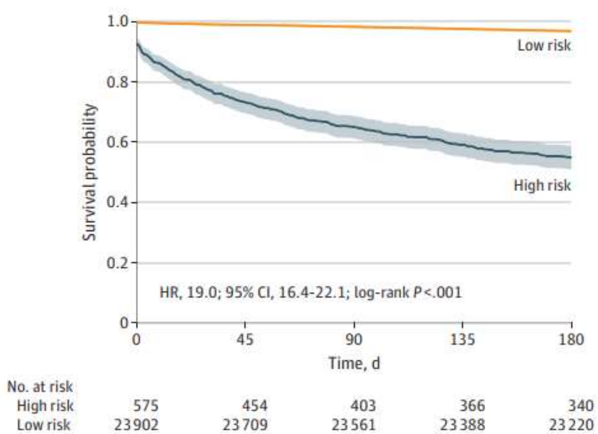

High-risk patients defined as having greater than 40% risk of 180-day mortality.

JAMA Oncology | Original Investigation

Effect of Integrating Machine Learning Mortality Estimates With Behavioral Nudges to Clinicians on Serious Illness Conversations Among Patients With Cancer
A Stepped-Wedge Cluster Randomized Clinical Trial

Christopher R. Manz, MD; Ravi B. Parikh, MD, MPP; Dylan S. Small, PhD; Chalanda N. Evans, BS; Corey Chivers, PhD; Susan H. Regli, PhD; C. William Hanson, MD; Justin E. Bekelman, MD; Charles A. L. Rareshide, MS; Nina O'Connor, MD; Lynn M. Schuchter, MD; Lawrence N. Shulman, MD; Mitesh S. Patel, MD, MBA

Table 2. Adjusted Changes in Serious Illness Conversations and in Advanced Care Planning

Conversation	No./total No. (%) of encounters		Adjusted difference for intervention relative to control, percentage points (95% CI)[a]	*P* value
	Control	Intervention		
Serious illness encounters				
All patients	155/12 170 (1.3)	632/13 889 (4.6)	3.3 (2.3-4.5)	<.001
High-risk patients	59/1732 (3.4)	246/1820 (13.5)	10.1 (6.9-13.8)	<.001
Advanced care planning encounters				
All patients	231/12 170 (1.9)	680/13 889 (4.9)	3.0 (2.1-4.1)	.001
High-risk patients	108/1732 (6.2)	286/1820 (15.7)	9.5 (5.8-13.5)	<.001

But can we go deeper?

Can we also use unstructured data (>80% of data*)

*Hyoun-Joong Kong Managing Unstructured Big Data in Healthcare System. Healthc Inform Res. 2019 Jan; 25(1): 1–2.cdoi: 10.4258/hir.2019.25.1.1

Survival prognosis with deep learning of structured variables AND clinical notes

- 1, 390, 032 provider notes
- 12 ,876 ,137 lab values (200 most common labs)
- 1 ,451, 740 vital signs
- 357, 981 diagnoses (500 most common codes)
- 1 ,162 ,164 procedures (500 most common codes)
- 1 ,834 ,477 medication orders (500 most common meds)

Automated Survival Prediction in Metastatic Cancer Patients Using High-Dimensional Electronic Medical Record Data

Michael F. Gensheimer, A. Solomon Henry, Douglas J. Wood, Trevor J. Hastie, Sonya Aggarwal, Sara A. Dudley, Pooja Pradhan, Imon Banerjee, Eunpi Cho, Kavitha Ramchandran, Erqi Pollom, Albert C. Koong, Daniel L. Rubin, Daniel T. Chang

JNCI J Natl Cancer Inst (2019) 111(6): djy178

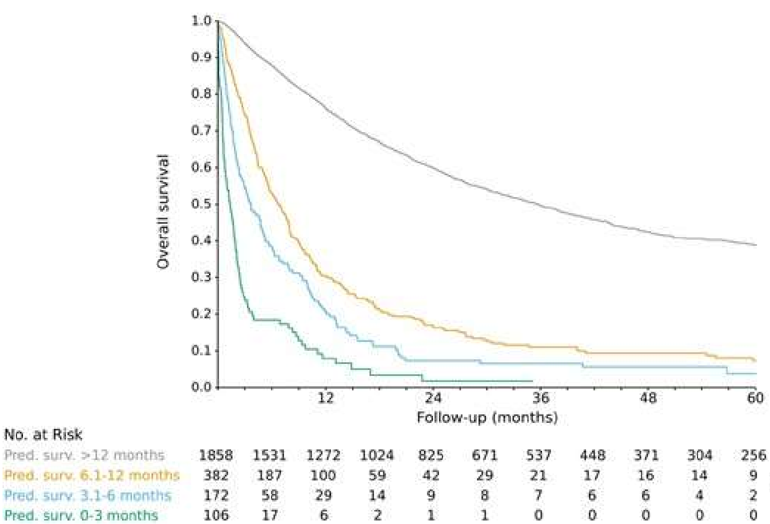

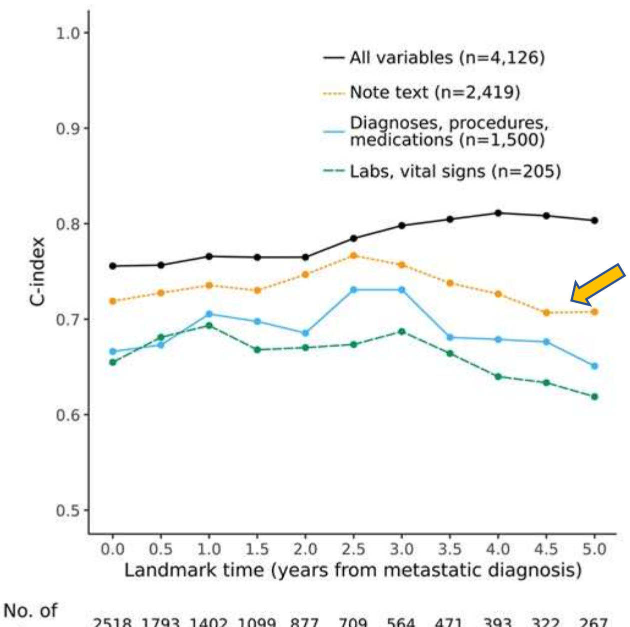

Table 2. Survival model coefficients for selected note text terms

Term	Coefficient*
Symptoms/appearance	
Cachectic	0.020
Fatigued	0.0059
Ascites	0.0085
Completely asymptomatic	−0.0054
Anxious	−0.0031
Feel well	−0.0073
Cancer location/response	
Disease progression	0.012
Leptomeningeal	0.0067
Mixed response	0.014
Innumerable pulmonary	0.0046
Minimal progression	−0.0012
Oligometastatic	−0.0066
Systemic therapy agents	
Nivolumab	−0.00065
Liposomal doxorubicin†	0.011
Anastrozole†	−0.00051
Leuprolide†	−0.0037
Tamoxifen	−0.0034

*A positive coefficient indicates shorter survival.
†Brand name converted to generic name for display.

But what if the outcome is not self-labeled?

Can we predict **ill-defined** events like cancer progression or treatment response?

ConvNet-based architecture on clinical text can detect cancer outcomes

JAMA Oncology | Original Investigation

Assessment of Deep Natural Language Processing in Ascertaining Oncologic Outcomes From Radiology Reports

Kenneth L. Kehl, MD, MPH; Haitham Elmarakeby, PhD; Mizuki Nishino, MD, MPH; Eliezer M. Van Allen, MD; Eva M. Lepisto, MA, MSc; Michael J. Hassett, MD, MPH; Bruce E. Johnson, MD; Deborah Schrag, MD, MPH

eFigure 1B: Deep learning architecture for ascertaining additional specific outcomes, including the presence of cancer, cancer progression, and cancer response, and specific sites of disease

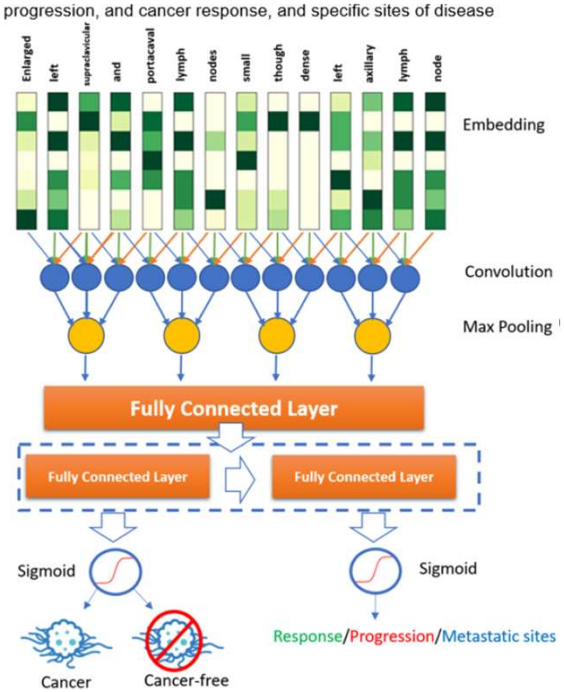

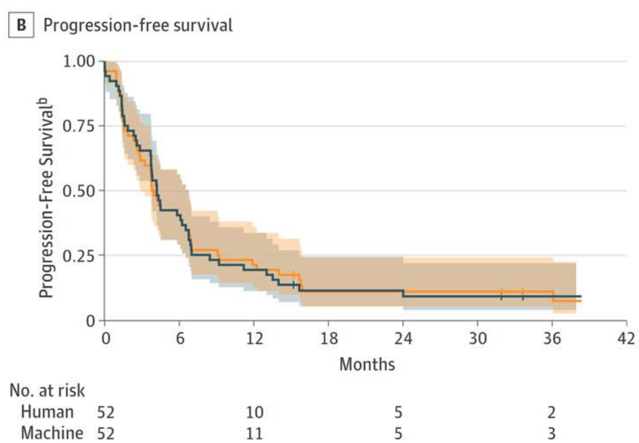

eFigure 3A: Example of prediction regarding any cancer

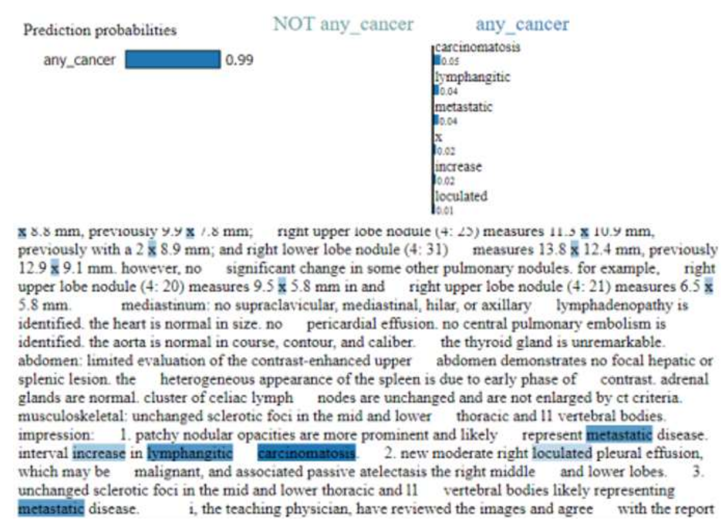

What if we wanted to understand the words in clinical notes?

Can semantic understanding of increase **performance** and/or improve **interpretability**?

A primer on word embeddings

"You shall know a word by the company it keeps"

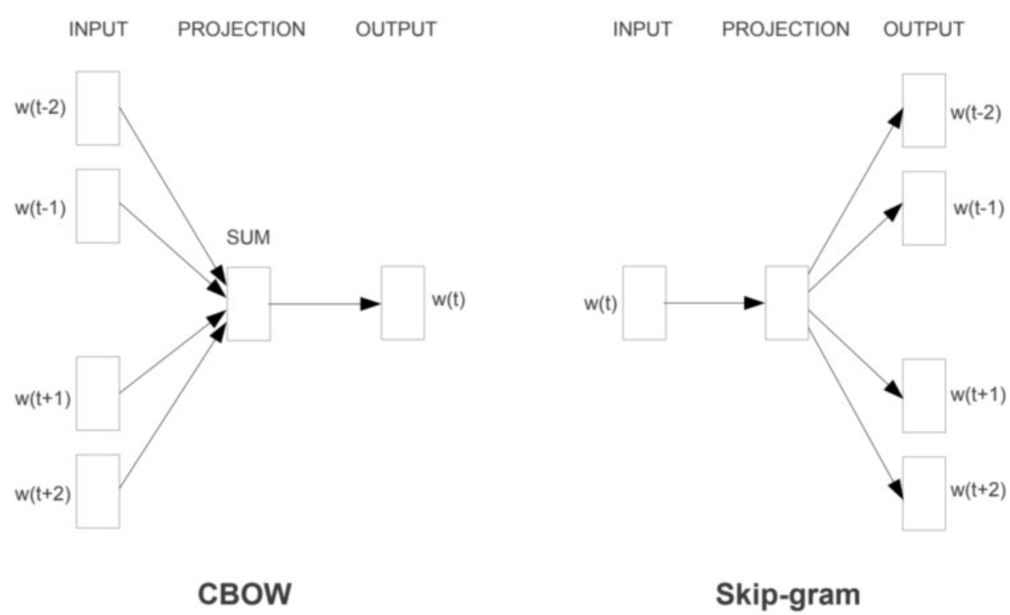

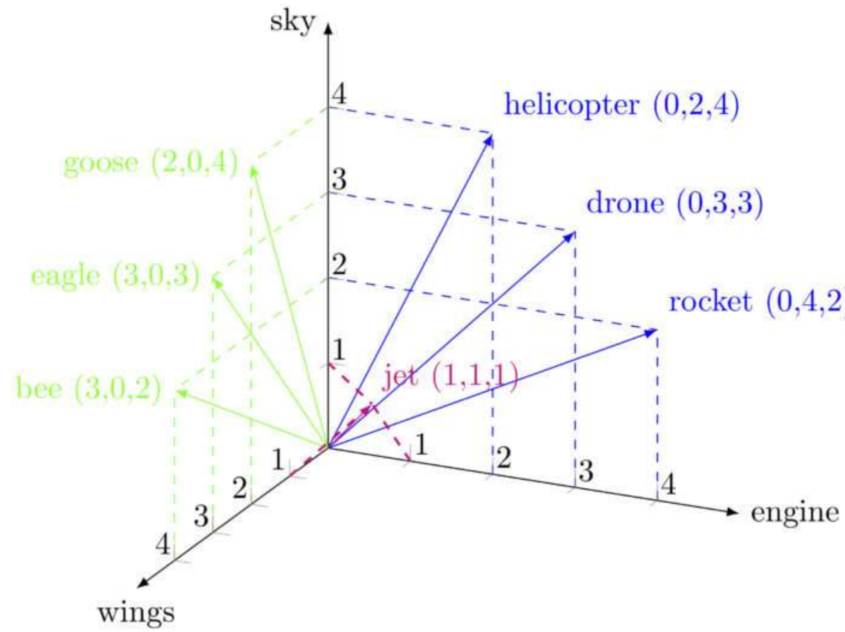

Mikolov et al. 2011 (Word2Vec)

Guillaume Desagulier tutorial
https://corpling.hypotheses.org/495

Combining semantic map with word embeddings increases interpretability

Radiology report annotation using intelligent word embeddings: Applied to multi-institutional chest CT cohort

Imon Banerjee[a,*], Matthew C. Chen[b], Matthew P. Lungren[b,*,1], Daniel L. Rubin[a,b,*,1]

[a] Department of Biomedical Data Science, Stanford University, Stanford, CA, United States
[b] Department of Radiology, Stanford University, Stanford, CA, United States

Journal of Biomedical Informatics 77 (2018) 11–20

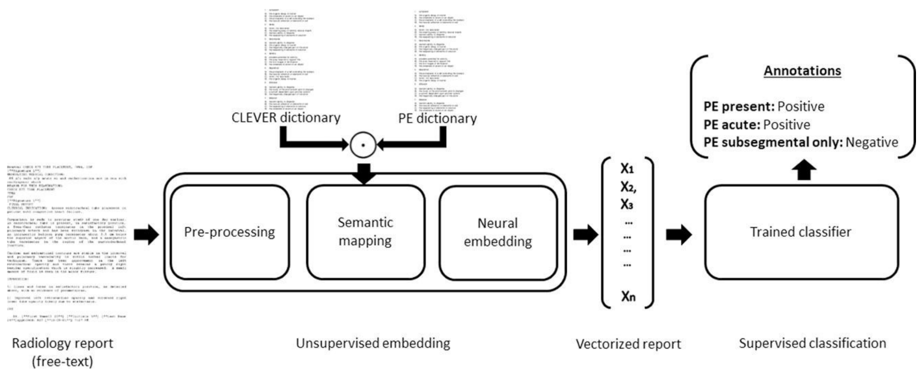

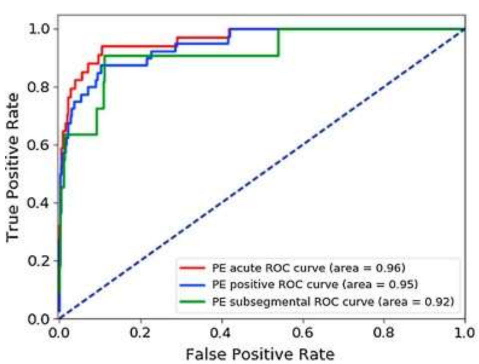

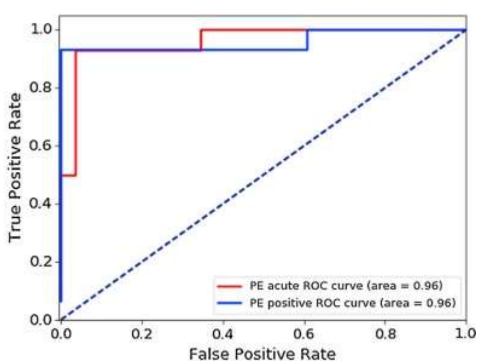

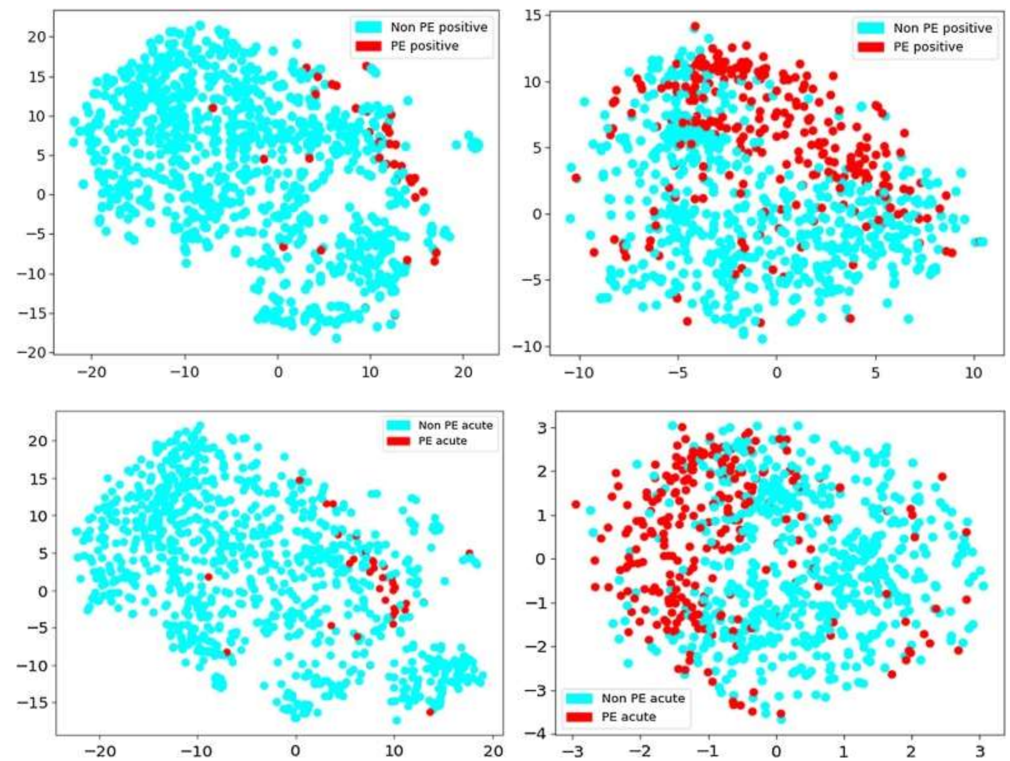

What if we do not even know what the labels are?

Can labels be extracted from unlabeled text?

Significant resources used to track radiation oncology research

2013

The Profession

National Institutes of Health Funding in Radiation Oncology: A Snapshot

Michael Steinberg, MD, William H. McBride, PhD, DSc, Erina Vlashi, PhD, and Frank Pajonk, MD, PhD

Department of Radiation Oncology, David Geffen School of Medicine at University of California, Los Angeles (UCLA), and Jonsson Comprehensive Cancer Center at UCLA, Los Angeles, California

Received Dec 20, 2012, and in revised form Jan 22, 2013. Accepted for publication Jan 27, 2013

"At the start of fiscal year 2013 we extracted records for 952 individual grants, which were active at the time of analysis from the NIH database…Our analysis identified 197 grants in radiation oncology. "

2014

Education Original Article

Current Status and Recommendations for the Future of Research, Teaching, and Testing in the Biological Sciences of Radiation Oncology: Report of the American Society for Radiation Oncology Cancer Biology/Radiation Biology Task Force, Executive Summary

Paul E. Wallner, DO,[*] Mitchell S. Anscher, MD,[†] Christopher A. Barker, MD,[‡] Michael Bassetti, MD, PhD,[§] Robert G. Bristow, MD, PhD,[‖] Yong I. Cha, MD, PhD,[¶] Adam P. Dicker, MD, PhD,[#] Silvia C. Formenti, MD,[**] Edward E. Graves, PhD,[††] Stephen M. Hahn, MD,[‡‡] Tom K. Hei, PhD,[§§] Alec C. Kimmelman, MD, PhD,[‖‖] David G. Kirsch, MD, PhD,[¶¶] Kevin R. Kozak, MD, PhD,[##] Theodore S. Lawrence, MD, PhD,[***] Brian Marples, PhD,[†††] William H. McBride, DSc,[‡‡‡] Ross B. Mikkelsen, PhD,[†] Catherine C. Park, MD,[§§§] Joanne B. Weidhaas, MD, PhD,[‖‖‖] Anthony L. Zietman, MD,[¶¶¶] and Michael Steinberg, PhD[‡‡‡]

"The first was…to congress about actual radiation oncology funding levels; the second was a review of the publicly available grant system database…To differentiate biological research from clinical trials and physics research, all radiation oncology grants….were hand-curated, separating the biology grants from the clinical and physics grants. Further, the biology grants were then subdivided by research topic."

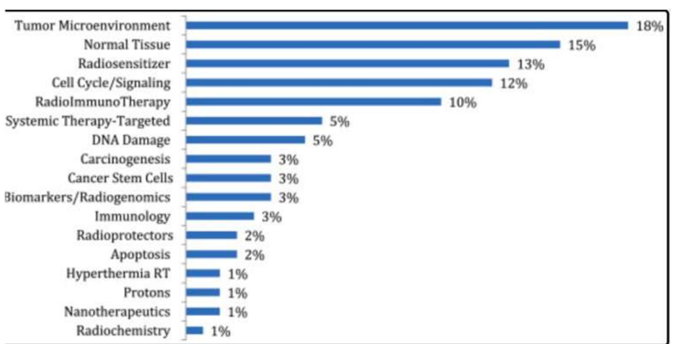

2017

The Profession

Analysis of the 2017 American Society for Radiation Oncology (ASTRO) Research Portfolio

James B. Yu, MD,[*] Tyler F. Beck, PhD,[†] Mitchell S. Anscher, MD,[‡] Andrew M. Baschnagel, MD,[§] Kristy K. Brock, PhD,[‖] David J. Carlson, PhD,[*] Michael M. Dominello, DO,[¶] Randall J. Kimple, MD, PhD,[§] Jonathan P. Knisely, MD,[#] Marc S. Mendonca, PhD,[**] Omar Y. Mian, MD, PhD,[††] Anurag K. Singh, MD,[‡‡] Eduardo G. Moros, PhD,[§§] and Judith C. Keen, PhD[†]

[*]Department of Radiation Oncology, Yale School of Medicine, New Haven, Connecticut; [†]American Society for Radiation Oncology, Arlington, Virginia; [‡]Department of Radiation Oncology, MD Anderson Cancer Center, Houston, Texas; [§]Department of Human Oncology, University of Wisconsin School of Medicine and Public Health, Madison, Wisconsin; [‖]Departments of Imaging Physics and Radiation Physics, MD Anderson Cancer Center, Houston, Texas; [¶]Department of Radiation Oncology, Karmanos Cancer Institute, Detroit, Michigan; [#]Department of Radiation Oncology, Weill Cornell Medicine, New York, New York; [**]Department of Radiation Oncology, Indiana University School of Medicine, Indianapolis, Indiana; [††]Department of Radiation Oncology, Cleveland Clinic, Cleveland, Ohio; [‡‡]Department of Radiation Oncology, Roswell Park Comprehensive Cancer Center, Buffalo, New York; and [§§]Department of Radiation Oncology, Moffitt Cancer Center, Tampa, Florida

Received Mar 22, 2018. Accepted for publication Jul 22, 2018.

Of the grants submitted…a significant number of grants were categorized as "unknown."

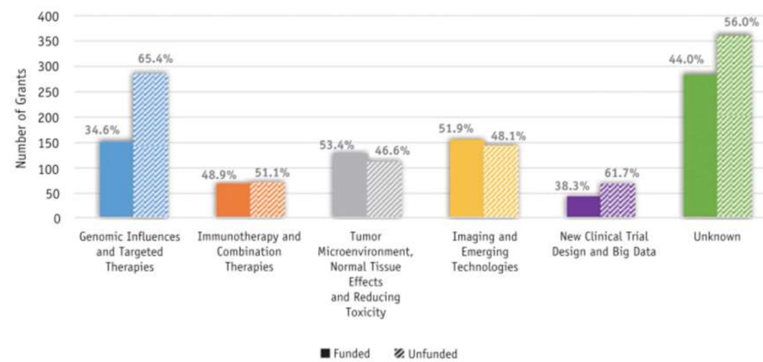

What ideas/topics/themes are being funded by NCI in Dept. of Radiology or Radiation Oncology?

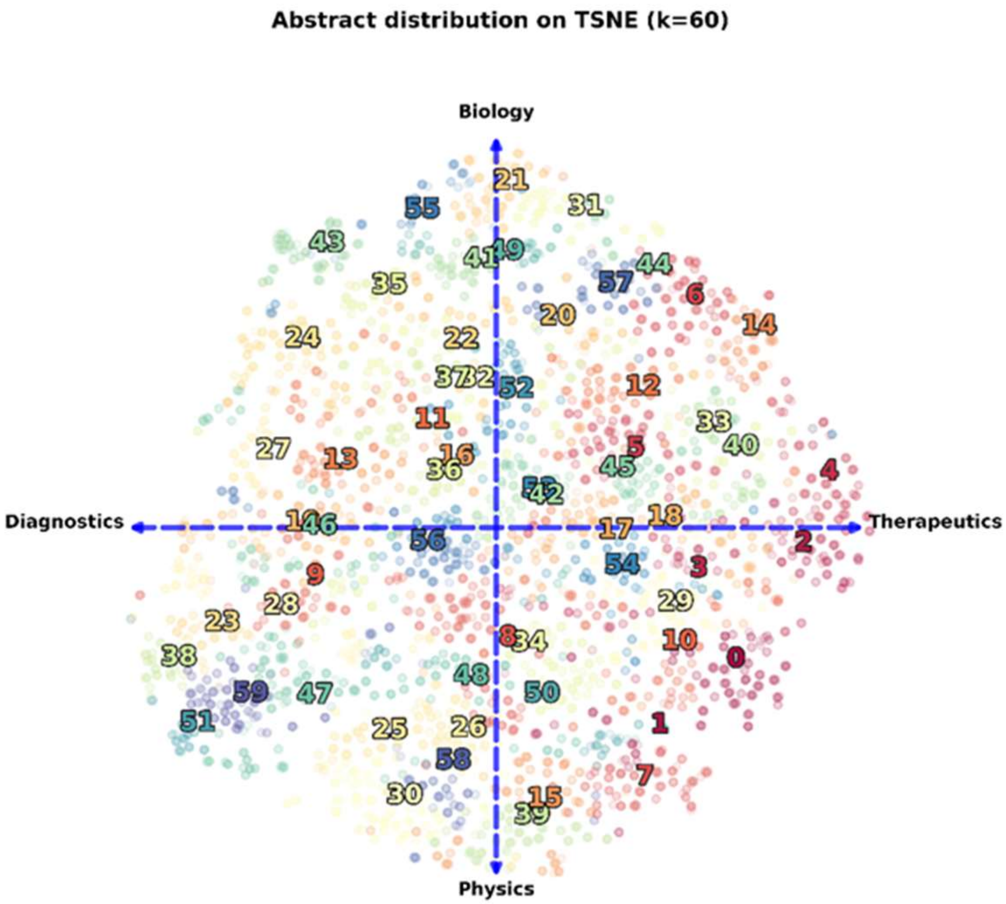

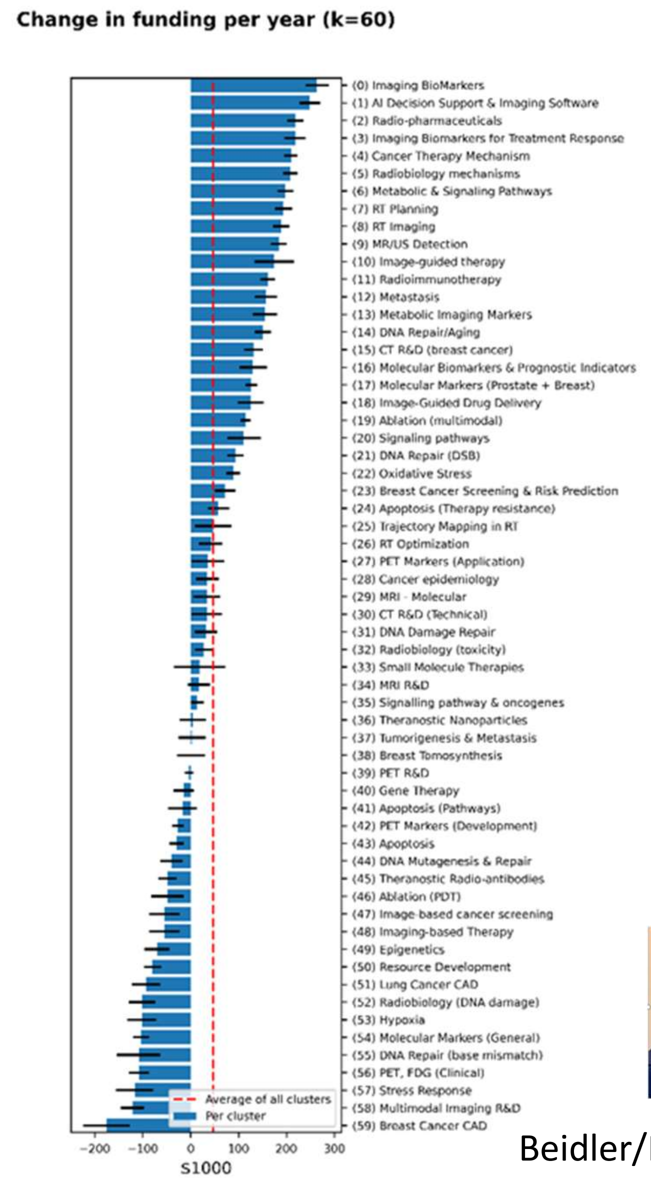

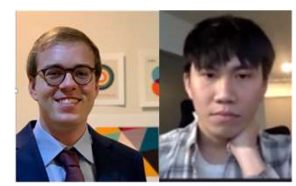

Beidler/Nguyen et al., in prep

Methods/results

-7k grant abstracts converted to BioWordVec embeddings (trained on biomedical+clinical data)

-clustered using combined hierarchical/K-means clustering

-used k=15 centroids (per elbow plot of clustering performance) and [k=60 (more realistic)]

-manual validated ~5% of grants over 4 raters (different training level) with good concordance

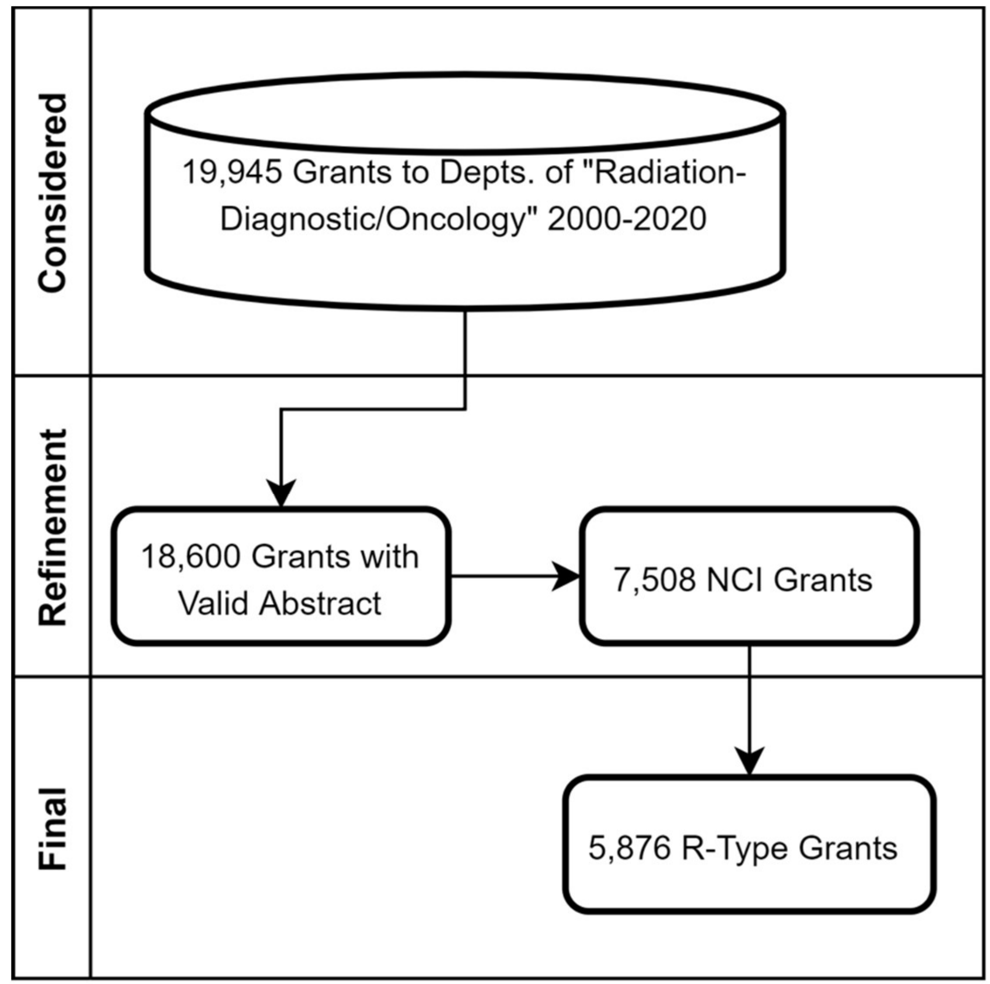

5-7k abstracts→15 domains

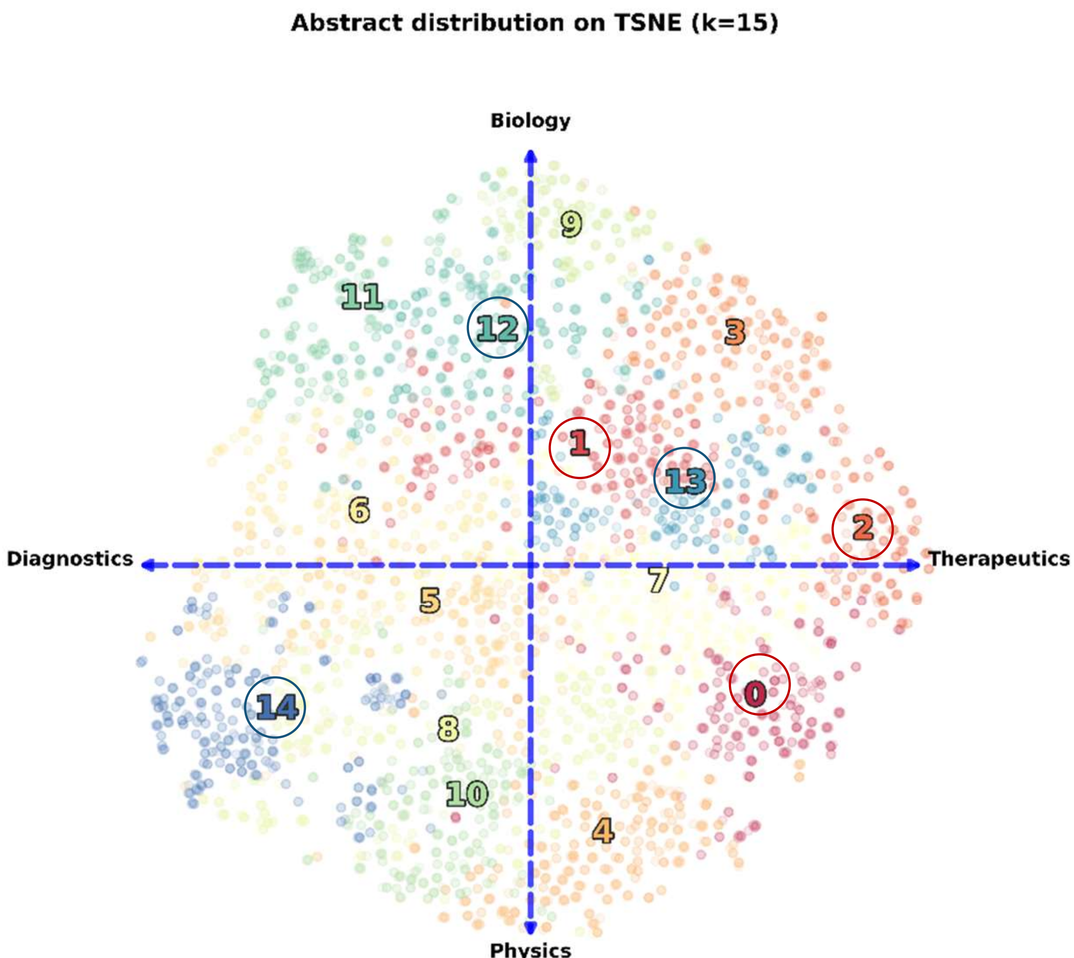

Beidler/Nguyen et al., in prep

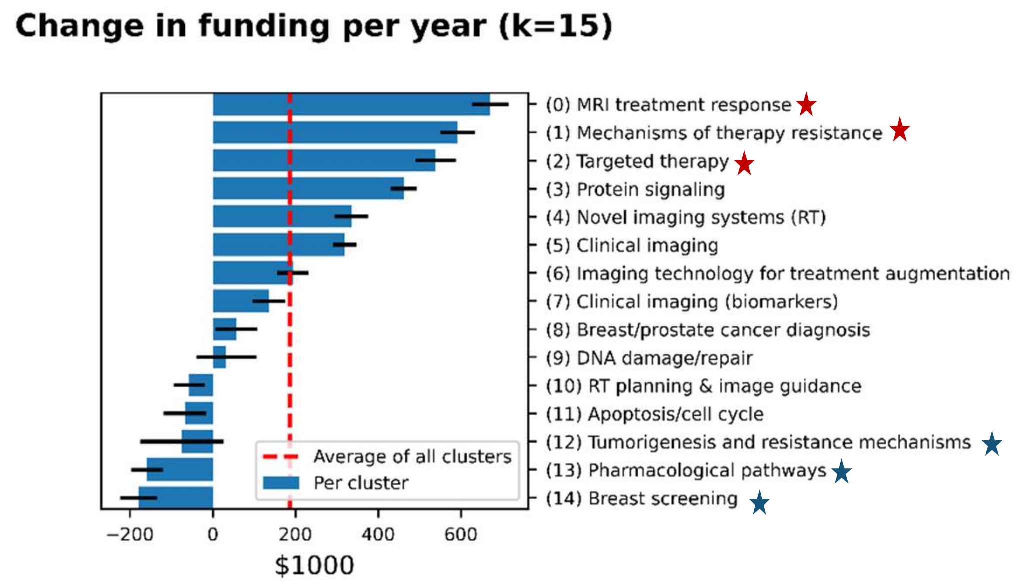

- 3 fastest growing (15 topics):
  - (0) MRI Treatment response
  - (1) Mechanisms of therapy resistance
  - (2) Targeted therapies
- 3 slowest growing (15 topics):
  - (14) Breast screening
  - (13) Pharmacological pathways
  - (12) Tumorigenesis and resistance mechanisms

5-7k abstracts->60 domains

- 3 fastest growing (60 topics):
  - (0) Imaging Biomarkers
  - (1) AI Decision Support & Imaging software
  - (2) Radio-pharmaceuticals
- 3 slowest growing (60 topics):
  - (59) Breast cancer CAD
  - (58) Multimodal Imaging R&D
  - (57) Stress Response

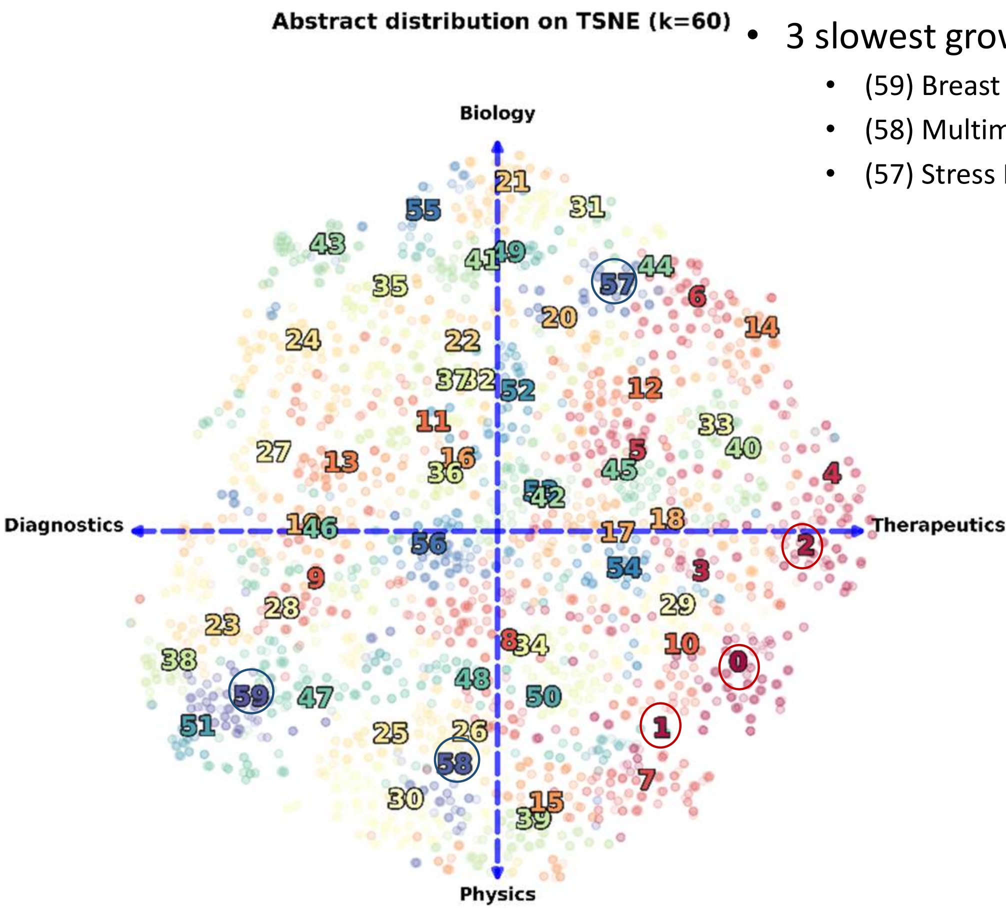

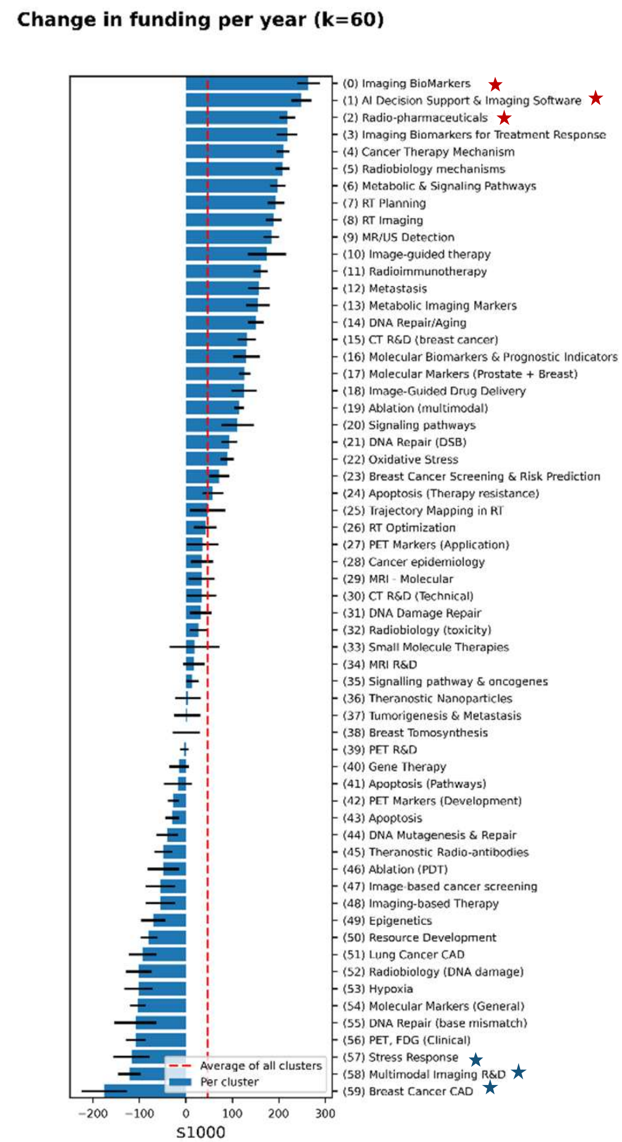

Manual validation shows reasonable concordance between human and machines

Not shown: experience level correlates with concordance

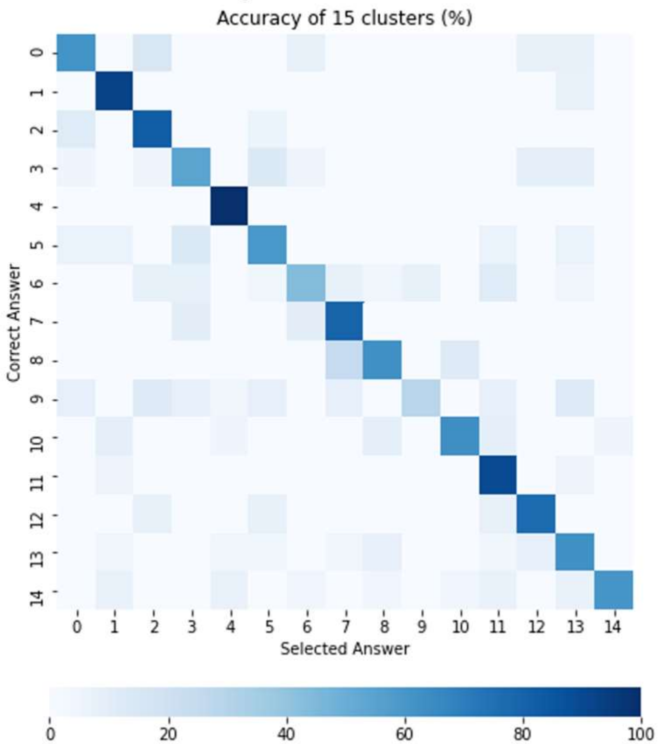

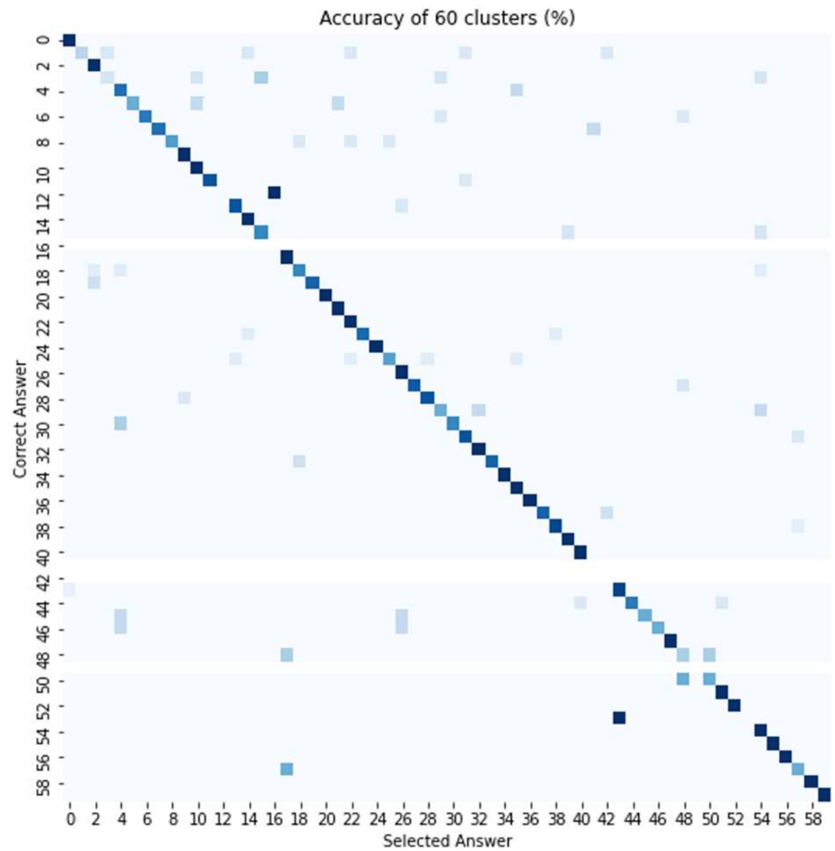

Grants wrongly labelled by the algorithm tended to be further from centroid

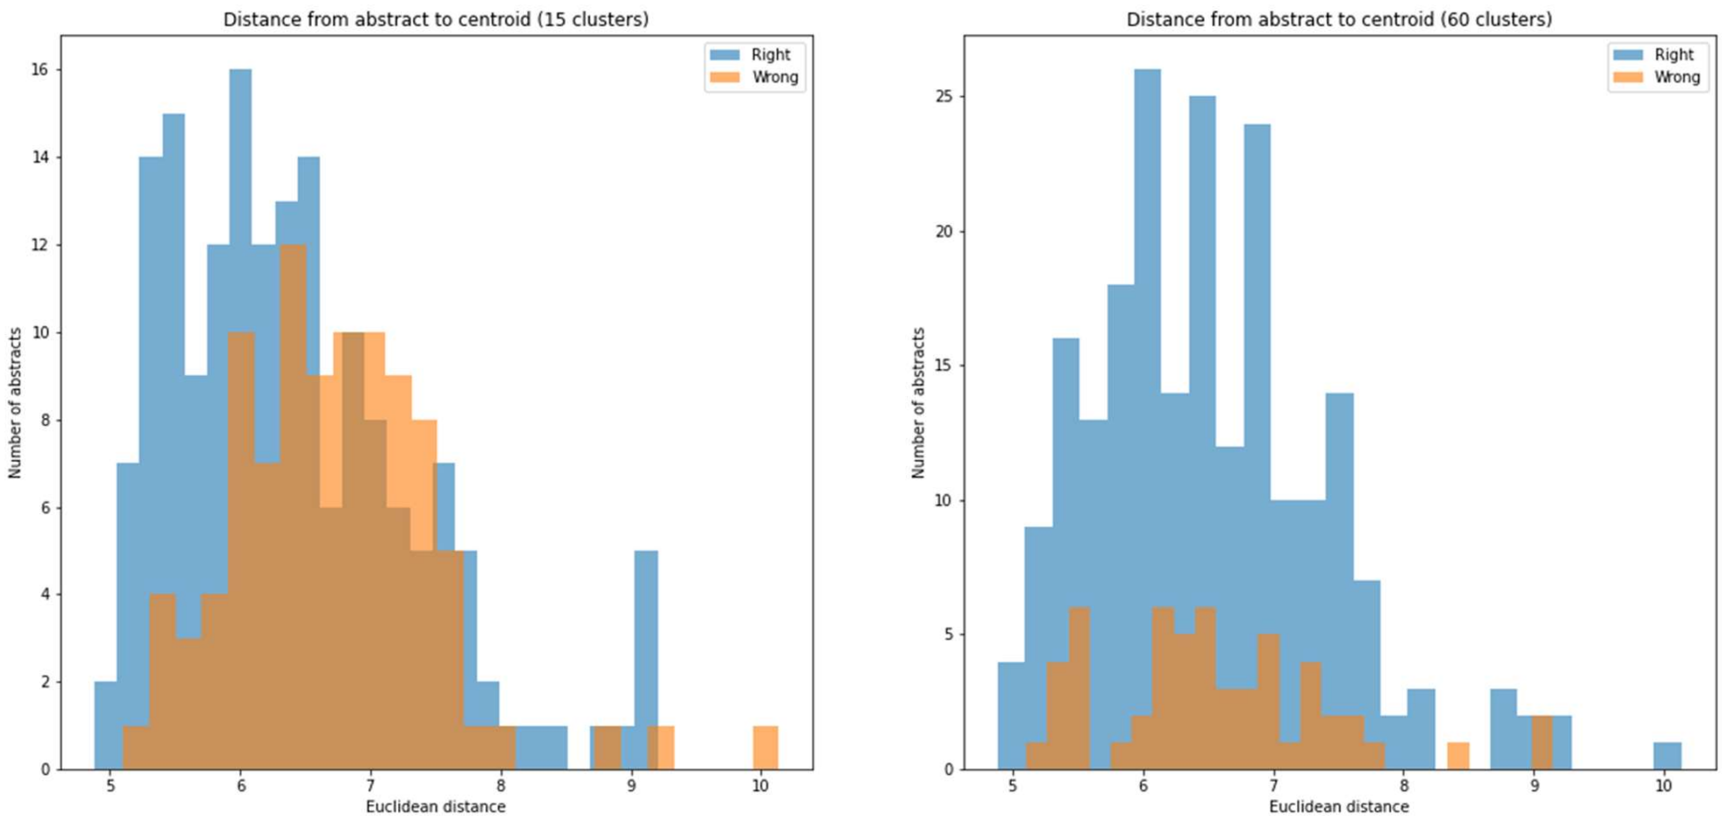

Funding topics have “emerged” and “disappeared” in last 20 years

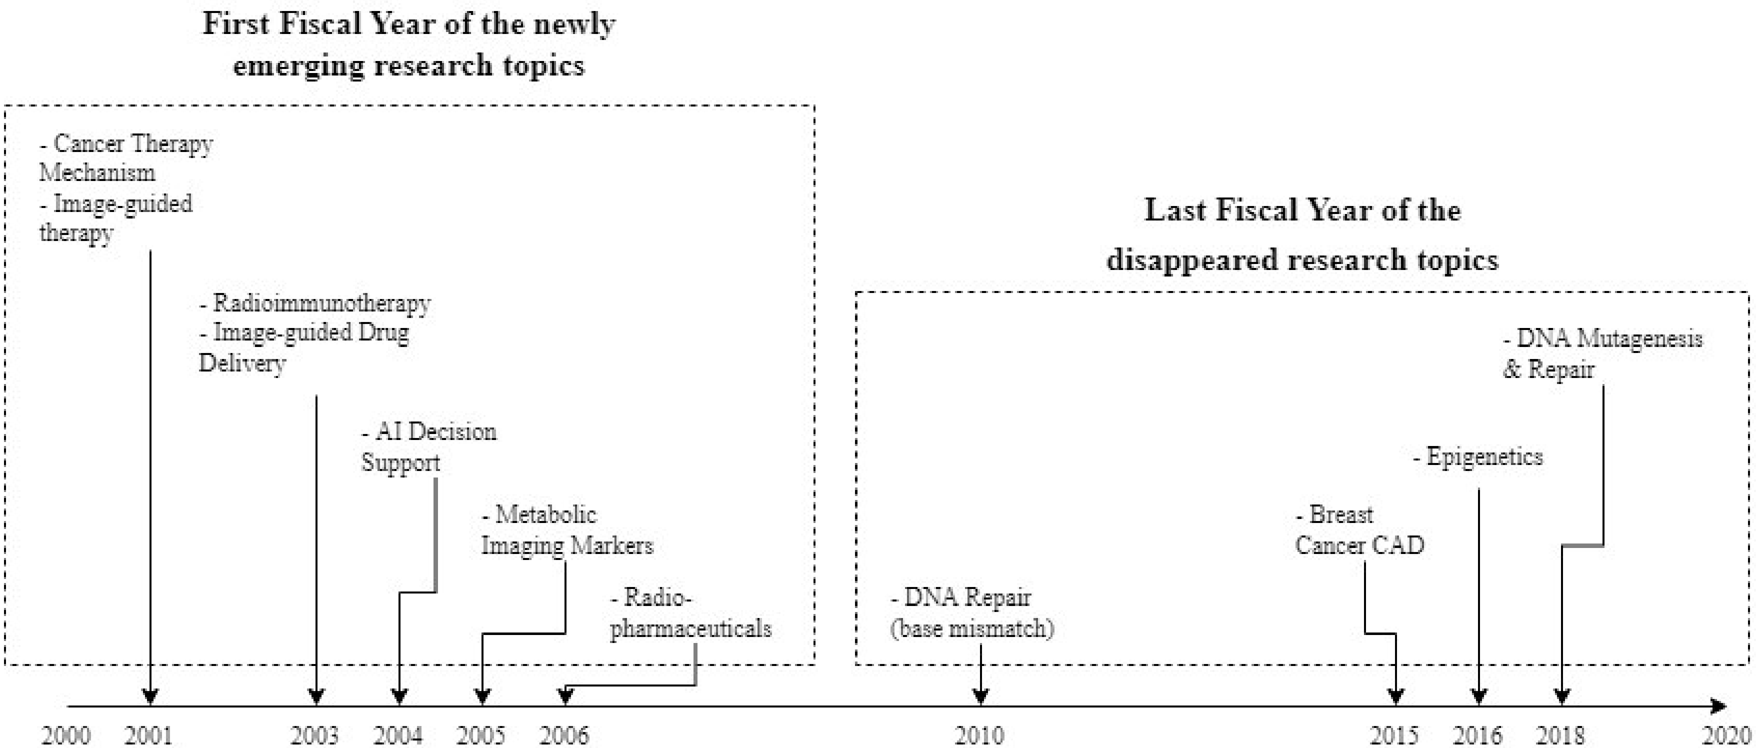

Limitations

- Grants further away from the centroid may not seem like they belong
- If new data is entered, the new optimal clustering may appear different
- 1 grant : 1 topic

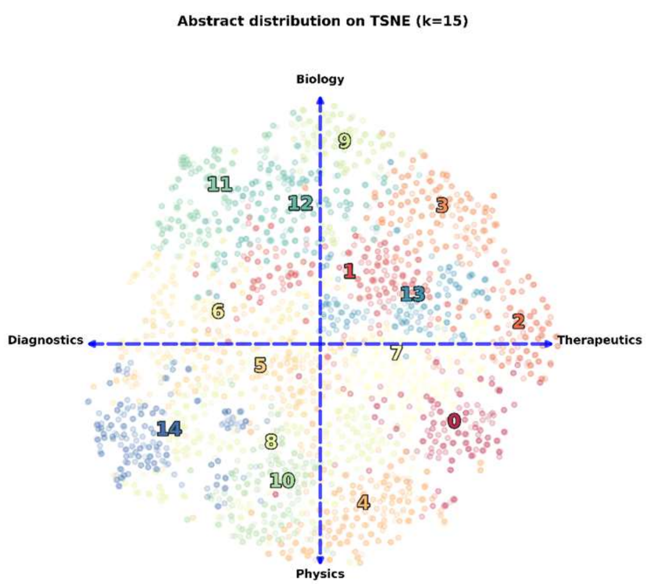

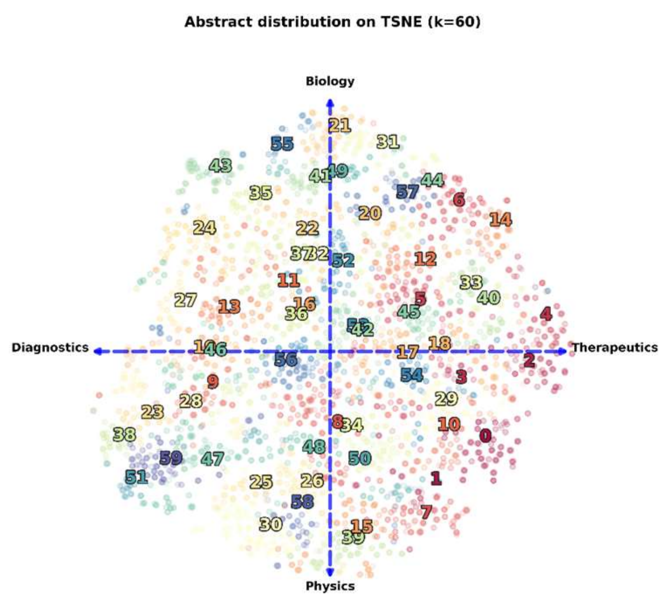

If a clinical model performs great but never affects patients, is it a useful model?

Mentees/trainees for this work

- August Anderson

- Peter Beidler

- Mark Nguyen

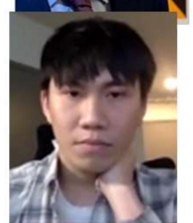

- Joseph Tsai

- Qian Zhang
  (UW→Northwestern)

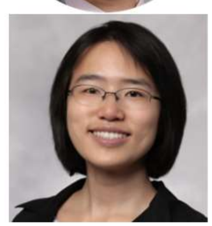

Collaborators for this work

- Eric Ford

- Kevin Lybarger
  (UW→George Mason
  https://www.kevinlybarger.me