INTERPRETABLE SEMANTIC VECTORS FROM A JOINT MODEL OF BRAIN- AND TEXT- BASED MEANING: SUPPLEMENTARY MATERIAL

ALONA FYSHE, PARTHA P TALUKDAR, BRIAN MURPHY, TOM M MITCHELL

1. Perceptual Features

These are the perceptual features used during the training of JNNSE to explain the non-semantic portion of the word/picture stimuli.

(1) Word length
(2) White pixel count
(3) Internal details
(4) Verticality
(5) Horizontalness
(6) Left-diagonalness
(7) Right-diagonalness
(8) Aspect-ratio: skinny?fat
(9) Prickiliness
(10) Line curviness
(11) 3D curviness

2. 218 Semantic Features

These semantic features were collected with Mechanical Turk and represent behavioral data for the 60 words for which we have brain activation data.

(1) Is it an animal?
(2) Is it a body part?
(3) Is it a building?
(4) Is it a building part?
(5) Is it clothing?
(6) Is it furniture?
(7) Is it an insect?
(8) Is it a kitchen item?
(9) Is it man-made?
(10) Is it a tool?
(11) Can you eat it?
(12) Is it a vehicle?
(13) Is it a person?

(14) Is it a vegetable/plant?
(15) Is it a fruit?
(16) Is it made of metal?
(17) Is it made of plastic?
(18) Is part of it made of glass?
(19) Is it made of wood?
(20) Is it shiny?
(21) Can you see through it?
(22) Is it colorful?
(23) Does it change color?
(24) Is one more than one colored?
(25) Is it always the same color(s)?
(26) Is it white?
(27) Is it red?
(28) Is it orange?
(29) Is it flesh-colored?
(30) Is it yellow?
(31) Is it green?
(32) Is it blue?
(33) Is it silver?
(34) Is it brown?
(35) Is it black?
(36) Is it curved?
(37) Is it straight?
(38) Is it flat?
(39) Does it have a front and a back?
(40) Does it have a flat/straight top?
(41) Does it have flat/straight sides?
(42) Is taller than it is wide/long?
(43) Is it long?
(44) Is it pointed/sharp?
(45) Is it tapered?
(46) Is it round?
(47) Does it have corners?
(48) Is it symmetrical?
(49) Is it hairy?
(50) Is it fuzzy?
(51) Is it clear?
(52) Is it smooth?
(53) Is it soft?
(54) Is it heavy?
(55) Is it lightweight?
(56) Is it dense?

(57) Is it slippery?
(58) Can it change shape?
(59) Can it bend?
(60) Can it stretch?
(61) Can it break?
(62) Is it fragile?
(63) Does it have parts?
(64) Does it have moving parts?
(65) Does it come in pairs?
(66) Does it come in a bunch/pack?
(67) Does it live in groups?
(68) Is it part of something larger?
(69) Does it contain something else?
(70) Does it have internal structure?
(71) Does it open?
(72) Is it hollow?
(73) Does it have a hard inside?
(74) Does it have a hard outer shell?
(75) Does it have at least one hole?
(76) Is it alive?
(77) Was it ever alive?
(78) Is it a specific gender?
(79) Is it manufactured?
(80) Was it invented?
(81) Was it around 100 years ago?
(82) Are there many varieties of it?
(83) Does it come in different sizes?
(84) Does it grow?
(85) Is it smaller than a golfball?
(86) Is it bigger than a loaf of bread?
(87) Is it bigger than a microwave oven?
(88) Is it bigger than a bed?
(89) Is it bigger than a car?
(90) Is it bigger than a house?
(91) Is it taller than a person?
(92) Does it have a tail?
(93) Does it have legs?
(94) Does it have four legs?
(95) Does it have feet?
(96) Does it have paws?
(97) Does it have claws?
(98) Does it have horns/thorns/spikes?
(99) Does it have hooves?

(100) Does it have a face?
(101) Does it have a backbone?
(102) Does it have wings?
(103) Does it have ears?
(104) Does it have roots?
(105) Does it have seeds?
(106) Does it have leaves?
(107) Does it come from a plant?
(108) Does it have feathers?
(109) Does it have some sort of nose?
(110) Does it have a hard nose/beak?
(111) Does it contain liquid?
(112) Does it have wires or a cord?
(113) Does it have writing on it?
(114) Does it have wheels?
(115) Does it make a sound?
(116) Does it make a nice sound?
(117) Does it make sound continuously when active?
(118) Is its job to make sounds?
(119) Does it roll?
(120) Can it run?
(121) Is it fast?
(122) Can it fly?
(123) Can it jump?
(124) Can it float?
(125) Can it swim?
(126) Can it dig?
(127) Can it climb trees?
(128) Can it cause you pain?
(129) Can it bite or sting?
(130) Does it stand on two legs?
(131) Is it wild?
(132) Is it a herbivore?
(133) Is it a predator?
(134) Is it warm blooded?
(135) Is it a mammal?
(136) Is it nocturnal?
(137) Does it lay eggs?
(138) Is it conscious?
(139) Does it have feelings?
(140) Is it smart?
(141) Is it mechanical?
(142) Is it electronic?

(143) Does it use electricity?
(144) Can it keep you dry?
(145) Does it provide protection?
(146) Does it provide shade?
(147) Does it cast a shadow?
(148) Do you see it daily?
(149) Is it helpful?
(150) Do you interact with it?
(151) Can you touch it?
(152) Would you avoid touching it?
(153) Can you hold it?
(154) Can you hold it in one hand?
(155) Do you hold it to use it?
(156) Can you play it?
(157) Can you play with it?
(158) Can you pet it?
(159) Can you use it?
(160) Do you use it daily?
(161) Can you use it up?
(162) Do you use it when cooking?
(163) Is it used to carry things?
(164) Can you pick it up?
(165) Can you control it?
(166) Can you sit on it?
(167) Can you ride on/in it?
(168) Is it used for transportation?
(169) Could you fit inside it?
(170) Is it used in sports?
(171) Do you wear it?
(172) Can it be washed?
(173) Is it cold?
(174) Is it cool?
(175) Is it warm?
(176) Is it hot?
(177) Is it unhealthy?
(178) Is it hard to catch?
(179) Can you peel it?
(180) Can you walk on it?
(181) Can you switch it on and off?
(182) Can it be easily moved?
(183) Do you drink from it?
(184) Does it go in your mouth?
(185) Is it tasty?

(186) Is it used during meals?
(187) Does it have a strong smell?
(188) Does it smell good?
(189) Does it smell bad?
(190) Is it usually inside?
(191) Is it usually outside?
(192) Would you find it on a farm?
(193) Would you find it in a school?
(194) Would you find it in a zoo?
(195) Would you find it in an office?
(196) Would you find it in a restaurant?
(197) Would you find in the bathroom?
(198) Would you find it in a house?
(199) Would you find it near a road?
(200) Would you find it in a dump/landfill?
(201) Would you find it in the forest?
(202) Would you find it in a garden?
(203) Would you find it in the sky?
(204) Do you find it in space?
(205) Does it live above ground?
(206) Does it get wet?
(207) Does it live in water?
(208) Can it live out of water?
(209) Do you take care of it?
(210) Does it make you happy?
(211) Do you love it?
(212) Would you miss it if it were gone?
(213) Is it scary?
(214) Is it dangerous?
(215) Is it friendly?
(216) Is it rare?
(217) Can you buy it?
(218) Is it valuable?

3. MEG Preprocessing

All subjects gave their written informed consent approved by the *anonymized* Institutional Review Board. MEG data were recorded using an Elekta Neuromag device (Elekta Oy). The data was acquired at 1 kHz, high-pass filtered at 0.1 Hz and low-pass filtered at 330 Hz. Eye movements (horizontal and vertical eye movements as well as blinks) were monitored by recording the differential activity of muscles above, below, and beside the eyes. At the beginning of each session we recorded the position of the subject's head with four head position indicator (HPI) coils placed on the subjects scalp. The HPI coils, along with three cardinal points (nasion, left and right pre-auricular), were digitized into the system.

The data were preprocessed using the Signal Space Separation method (SSS) [5, 3] and temporal extension of SSS (tSSS) [4] to remove artifacts and noise unrelated to brain activity. In addition, we used tSSS to realign the head position measured at the beginning of each block to a common location. The MEG signal was then low-pass filtered to 50 Hz to remove the contributions of line noise and down-sampled to 200 Hz. The Signal Space Projection method (SSP) [6] was then used to remove signal contamination by eye blinks or movements, as well as MEG sensor malfunctions or other artifacts. MEG recordings are known to drift with time, so we corrected our data by subtracting the mean signal amplitude during the 200ms before stimulus onset, for each sensor/repetition pair. Because the magnitude of the MEG signal is very small, we multiplied the signal by 10^{12} to avoid numerical precision problems. We used the 750 ms of MEG signal beginning immediately after the onset of the noun stimulus, which is the generally agreed upon time during which semantic processing occurs [2].

4. FMRI Preprocessing

The fMRI was collected by the authors of [1] and the data is available online (`http://www.cs.cmu.edu/~tom/science2008/`). Thus, the processing steps are exactly as described in the corresponding publication.

5. Dropout Test

To test that the brain activation data is truly influencing rows of A not constrained by brain activation data, we performed a *dropout* test. Again, we split the original 60 words into two 30 word groups (as evenly as possible across word categories). We trained JNNSE(fMRI+Text) with 30 words, simulating the scenario where we have corpus data but no brain data for some words and some subjects. We then use the rows of the resulting A matrix corresponding to the 30 withheld words to test 2 vs. 2 accuracy with the remaining 8 fMRI subjects. The testing data has been halved, so we used 75 randomly chosen word pairs instead of 150. Because the words are disjoint, we could have tested with all 9 subjects, but for the most accurate comparison, we followed the same methodology as the previous 2 vs. 2 test.

Along with the results from the original 2 vs. 2 results for all 60 words, Figure 1 shows the results for the dropout scenario for JNNSE(fMRI+Text) with $\ell = 1000$ latent dimensions tested on fMRI data. In this scenario, each time we perform a 2 vs. 2 test we are training on 28 instead of 58 words, and so we expect performance to suffer. For NNSE(Text), performance on the 2 vs. 2 task with only 30 words is very low, 55.6%. The drop in accuracy is due only to the reduction in training data, as there is no change in the semantic vectors used to perform the 2 vs. 2 test. The results are lower for JNNSE(fMRI+Text) tested on 30 words, but is still 7% higher than results with NNSE(Text). Because the training and testing words are completely disjoint, these results imply that the addition of brain activation data improves the learned latent representation, not only for those words for which we have brain activation data, but also for the words for which there is no brain activation data. This, along with the fact that all latent dimensions used by words with brain data are also used by words without brain data, suggests that our algorithm produces semantic representations that are better constrained for all words in A, even though we only add explicit additional constraints to a small number of words. The dropout test also shows that we could have collected a different set of 60 word for each of the 9 subjects for a total of 540 words and still successfully used JNNSE(Brain+Text). In the future, this realization will allow us to increase our coverage over words and could lead to greatly improved VSMs.

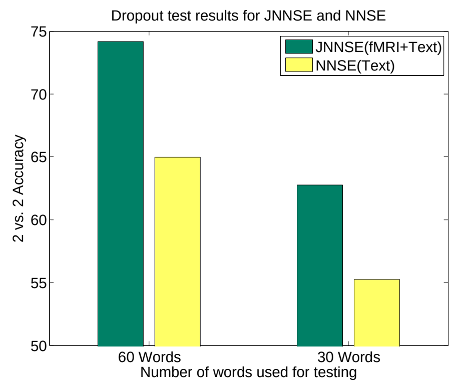

FIGURE 1. **Performance on the dropout test, as tested on fMRI.** JNNSE(fMRI+Text) with $\ell = 1000$ when train/tested on the same 60 words (same as Figure 2 of main text) and when train/tested on disjoint sets of 30 words. Performance decreases for both JNNSE(fMRI+Text) and NNSE(Text), but JNNSE(fMRI+Text) still outperforms NNSE(Text).

REFERENCES

[1] Tom M Mitchell, Svetlana V Shinkareva, Andrew Carlson, Kai-Min Chang, Vicente L Malave, Robert A Mason, and Marcel Adam Just. Predicting human brain activity associated with the meanings of nouns. *Science (New York, N.Y.)*, 320(5880):1191–5, May 2008.

[2] Riitta Salmelin. Clinical neurophysiology of language: the MEG approach. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*, 118(2):237–54, February 2007.

[3] S Taulu and J Simola. Spatiotemporal signal space separation method for rejecting nearby interference in MEG measurements. *Physics in Medicine and Biology*, 51:1–10, 2006.

[4] Samu Taulu and Riitta Hari. Removal of magnetoencephalographic artifacts with temporal signal-space separation: demonstration with single-trial auditory-evoked responses. *Human brain mapping*, 30(5):1524–34, May 2009.

[5] Samu Taulu, Matti Kajola, and Juha Simola. The Signal Space Separation method. *ArXiv Physics*, 2004.

[6] M A Uusitalo and R J Ilmoniemi. Signal-space projection method for separating MEG or EEG into components. *Medical & biological engineering & computing*, 35(2):135–40, March 1997.