, the free encyclopedia that anyone can edit. 107,766 active editors 7,028,184 articles in English Lesley James McNair (25 May 1883 – 25 July 1944) was a lieutenant general of the United States Army who served in both world wars, and previously saw service in the Veracruz occupation and the Pancho Villa Expedition. During World War I, he served with the American Expeditionary Forces on the Western Front. At 35, he became the Army's second-youngest general officer. During the early stages of World War II, he was the commander of Army Ground Forces, and played the leading role in the organization, equipping, and training of Army units before they departed for overseas combat. He was killed on 25 July 1944 while in France as commander of the fictitious First United States Army Group, part of Operation Quicksilver, a deception plan for the invasion of Normandy. McNair died when the US Eighth Air Force attempted to use heavy bombers in support of ground combat troops, and several planes dropped payloads short of their targets. He received a posthumous promotion to general. (Full article...) Recently featured: Second Test, 1948 Ashes series Daily News Building Hippocampus Archive By email More featured articles About State flag of Transnistria ... that the state flag of Transnistria (pictured) features communist symbols, even though Transnistria is not a communist state? ... that actor Suja Khondokar served as a flight attendant aboard the inaugural commercial flight of Biman Bangladesh Airlines? ... that Peru's fishing industry is the world's largest producer of fishmeal and fish oil? ... that the Albanian noble Blasius Mataranga may have been captured and killed by the Balsha family—after they lured him into peace talks with a false promise of safe conduct? ... that Mayu Sakai found drawing Peter Pan Syndrome difficult because of its fantasy elements? ... that American football player Bob Wicks was teammates with the same quarterback from fourth grade through college? ... that from 1986 to 1988, a Boston TV station was the only one in the U.S. to broadcast digital audio in its video signal? ... that Archbishop Letard II of Nazareth was one of only a few fellow bishops about whom William of Tyre had something nice to say? ... that some Italians eat eggplants with chocolate? Archive Start a new article Nominate an article Ozzy Osbourne Ozzy Osbourne (pictured), the lead singer of Black Sabbath, dies at the age of 76. A fighter jet crashes into a college in Dhaka, Bangladesh, killing more than 30 people. In golf, Scottie Scheffler wins the Open Championship. A tourist boat capsizes during a thunderstorm in Hạ Long Bay, Vietnam, leaving at least 36 people dead. Ongoing: Gaza war Russian invasion of Ukraine timeline Sudanese civil war timeline Recent deaths: Hulk Hogan Thomas Anthony Durkin Giora Epstein Béatrice Uria-Monzon Rex White Roger Norrington Nominate an article July 25: National Day of Galicia, Saint James's Day, Tenjin Matsuri The Concorde involved in the Air France Flight 4590 accident 1261 – Alexios Strategopoulos led Nicaean forces to recapture Constantinople, leading to the reestablishment of the Byzantine Empire and the end of the Latin Empire. 1893 – The Corinth Canal was formally opened, bisecting the narrow Isthmus of Corinth in Greece to connect the Ionian Sea's Gulf of Corinth with the Aegean Sea's Saronic Gulf. 1950 – Korean War: After American troops withdrew, North Korean forces captured the village of Yongdong in South Korea. 2000 – Air France Flight 4590 (plane used pictured), a Concorde en route from Paris to New York, crashed in Gonesse, France, killing all 109 passengers on board and four people on the ground. 2010 – WikiLeaks published 75,000 classified documents about the War in Afghanistan in one of the largest leaks in U.S. military history. Sibylla, Queen of Jerusalem (d. 1190)Matt LeBlanc (b. 1967)Meg Donnelly (b. 2000)Azimzhan Askarov (d. 2020) More anniversaries: July 24 July 25 July 26 Archive By email List of days of the year About The Carousel at Glen Echo Park Amusement rides on the National Register of Historic Places (NRHP) are located throughout the United States. These individual ride listings consist mainly of carousels, but also include roller coasters, trains, and other ride types. Many NRHP-listed rides operate within amusement parks, with more than one present in Cedar Point, Lagoon, and Santa Cruz Beach Boardwalk. There are also high concentrations in New York City, the Greater Binghamton area in New York state, and Portland, Oregon. The first NRHP amusement ride listing was added in 1975 for the Idora Park Merry-Go-Round (delisted in 1985). The listing for the Crescent Park Looff Carousel was added in 1976 and is the ride that has existed on the NRHP the longest. The listing for the Portland Zoo Railway Historic District was created in 2020 and is the newest ride entry on the NRHP. Several NRHP-listed rides, including the Leap-the-Dips roller coaster, have a higher National Historic Landmark status. Of the nearly 100,000 NRHP listings, fewer than 100 are for amusement rides. (Full list...) Recently featured: One Direction discography Kerivoulines Accolades received by Inception Archive More featured lists Hudson Yards is a 28-acre (11-hectare) real-estate development located in Hudson Yards, a neighborhood of the New York City borough of Manhattan. It is situated on the waterfront of the Hudson River, on a platform built over the West Side Yard, a storage depot for the Long Island Rail Road. Related Companies and Oxford Properties are the primary developers and major equity partners in the project, with the master plan designed by the architectural firm Kohn Pedersen Fox. Construction began in 2012 and the first phase opened in 2019, with completion of the second phase expected by 2032. Major office tenants in the development include Warner Bros. Discovery, L'Oréal, and Wells Fargo among others. This photograph shows the skyscrapers of Hudson Yards, viewed across the Hudson River from Weehawken, New Jersey, in 2021. Photograph credit: Tony Jin Recently featured: Emperor angelfish Amália Rodrigues Atari video game burial Archive More featured pictures Community portal – The central hub for editors, with resources, links, tasks, and announcements. Village pump – Forum for discussions about Wikipedia itself, including policies and technical issues. Site news – Sources of news about Wikipedia and the broader Wikimedia movement. Teahouse – Ask basic questions about using or editing Wikipedia. Help desk – Ask questions about using or editing Wikipedia. Reference desk – Ask research questions about encyclopedic topics. Content portals – A unique way to navigate the encyclopedia. Wikipedia is written by volunteer editors and hosted by the Wikimedia Foundation, a non-profit organization that also hosts a range of other volunteer projects: CommonsFree media repository MediaWikiWiki software development Meta-WikiWikimedia project coordination WikibooksFree textbooks and manuals WikidataFree knowledge base WikinewsFree-content news WikiquoteCollection of quotations WikisourceFree-content library WikispeciesDirectory of species WikiversityFree learning tools WikivoyageFree travel guide WiktionaryDictionary and thesaurus This Wikipedia is written in English. Many other Wikipedias are available; some of the largest are listed below. 1,000,000+ articles العربية Deutsch Español فارسی Français Italiano Nederlands 日本語 Polski Português Русский Svenska Українська Tiếng Việt 中文 250,000+ articles Bahasa Indonesia Bahasa Melayu Bân-lâm-gú Български Català Čeština Dansk Eesti Ελληνικά Esperanto Euskara עברית Հայերեն Magyar Norsk bokmål Română Simple English Slovenčina Srpski Srpskohrvatski Suomi Türkçe Oʻzbekcha 50,000+ articles Asturianu Azərbaycanca Bosanski کوردی Frysk Gaeilge Galego Hrvatski ქართული Kurdî Latviešu Lietuvių Македонски Norsk nynorsk Shqip Slovenščina اردو Retrieved from " 2 Calendar year Years Millennium 2nd millennium Centuries 12th century 13th century 14th century Decades 1240s 1250s 1260s 1270s 1280s Years 1258 1259 1260 1261 1262 1263 1264 vte 1261 by topic Leaders Political entities State leaders Religious leaders Birth and death categories Births – Deaths Establishments and disestablishments categories Establishments – Disestablishments Art and literature 1261 in poetry vte 1261 in various calendarsGregorian calendar1261MCCLXIAb urbe condita2014Armenian calendar710ԹՎ ՉԺAssyrian calendar6011Balinese saka calendar1182–1183Bengali calendar667–668Berber calendar2211English Regnal year45 Hen. 3 – 46 Hen. 3Buddhist calendar1805Burmese calendar623Byzantine calendar6769–6770Chinese calendar庚申年 (Metal Monkey)3958 or 3751 — to — 辛酉年 (Metal Rooster)3959 or 3752Coptic calendar977–978Discordian calendar2427Ethiopian calendar1253–1254Hebrew calendar5021–5022Hindu calendars - Vikram Samvat1317–1318 - Shaka Samvat1182–1183 - Kali Yuga4361–4362Holocene calendar11261Igbo calendar261–262Iranian calendar639–640Islamic calendar659–660Japanese calendarBun'ō 2 / Kōchō 1(弘長元年)Javanese calendar1170–1171Julian calendar1261MCCLXIKorean calendar3594Minguo calendar651 before ROC民前651年Nanakshahi calendar−207Thai solar calendar1803–1804Tibetan calendar(male Iron-Monkey)1387 or 1006 or 234 — to — (female Iron-Bird)1388 or 1007 or 235 Michael VIII Palaiologos (1223–1282) Year 1261 (MCCLXI) was a common year starting on Saturday of the Julian calendar. March 13 – Treaty of Nymphaeum: Emperor Michael VIII Palaiologos signs a trade and defense agreement with the Republic of Genoa, to counterweight the Venetian presence in the region. Genoa agrees to ally with the Empire of Nicaea, by providing a fleet of up to 50 galleys during the projected Nicaean siege of Constantinople, while 16 galleys are to be immediately sent against the Latin Empire.[1] July – Michael sends his general Alexios Strategopoulos with a small advance force of 800 soldiers, most of them Cumans, to keep watch on the Bulgarians and scout the defending positions of the Latin forces in the surroundings of Constantinople. When they reach the village of Selymbria, Strategopoulos is informed by local farmers that the entire Latin garrison and the Venetian fleet, are absent conducting a raid against the Nicaean island of Daphnousia. He decides not to lose such a golden opportunity and makes plans (without the consent of Michael) to retake the capital.[2] July 25 – Reconquest of Constantinople: Alexios Strategopoulos and his men hide at a monastery near the city gates, before entering through a secret passage. After a short struggle, the guards who are completely taken by surprise are killed and the Venetian quarter is set ablaze. Panic spreads through the capital and Emperor Baldwin II rushes out to save his life, evacuating along with many other Latins with the help of the Venetian fleet. Baldwin manages to escape to the still Latin-held parts of Greece, but Constantinople is lost for good.[3] August 15 – Michael enters Constantinople in triumph and is crowned as emperor of the Byzantine Empire at the Hagia Sophia. To solidify his claim, the legitimate ruler, John IV Laskaris, is blinded on Michael's orders on December 25, his 11th birthday. Michael banishes him to a monastery and marries his two sisters to lesser Latin and Bulgarian nobles in an attempt to wipe out the Laskarid dynasty.[4] Kublai Khan releases 75 Chinese merchants who were captured along the border of the Mongol Empire. By doing this, Kublai hopes to bolster his popularity and depend on the cooperation of his Chinese subjects to ensure that his army receives more resources.[5] June 13 – Al-Mustansir II becomes the first Abbasid ruler in Cairo (after his escape during the Siege of Baghdad). He is sent with an army by Sultan Baibars to recover Baghdad, but is killed in a Mongol ambush near Anbar (modern Iraq), on November 28. The Abbasid caliphs continue as religious figureheads for the Mamluks in Egypt until the 16th century.[6] June 12 – King Henry III of England is released from his oath to maintain the Provisions of Oxford. He hires an army of 300 French knights as a bodyguard and takes up position in the Tower of London. He dismisses the baronial officials (led by Simon de Montfort) who wish the royal power to be modified by the principle of representation. This sets the stage for the Second Barons' War.[7] August – Battle of Callann in Ireland: Norman forces under John FitzThomas are defeated by a Gaelic army led by King Fínghin Mac Carthaigh. John FitzGerald is killed during the fighting.[8] February – Baldwin II and his son Philip of Courtenay ... reign of the 11-year-old Emperor Kameyama (until 1264). Early – Following disputes, northern academics from the University of Cambridge in England set up a University of Northampton by royal charter but it is suppressed by the Crown in 1265.[9] The earliest extant Chinese illustration of "Pascal's Triangle" is from Yang Hui's (or Qianguang's) book Xiangjie Jiuzhang Suanfa, published this year. May 25 – Pope Alexander IV dies after a pontificate of 6-years at Viterbo. He is succeeded by Urban IV as the 182nd pope of the Catholic Church. August 29 – Urban IV offers the crown of Sicily to Charles of Anjou, youngest son of King Louis VIII of France, hoping to strengthen his position. Wurmsbach Abbey (located in Bollingen) is established by Count Rudolf V of Rapperswil in Switzerland. February 1 – Walter de Stapledon, English bishop of Exeter (d. 1326) February 15 – Stephen I, duke of Bavaria, king of Hungary and Croatia (d. 1312) February 28 – Margaret of Scotland, queen consort of Norway (d. 1283) March 1 – Hugh Despenser the Elder, English chief adviser (d. 1326) July 25 – Arthur II, Breton nobleman (House of Dreux) (d. 1312) October 9 – Denis I ("the Poet King"), king of Portugal (d. 1325) Abu Abdallah ibn al-Hakim, Andalusian vizier and poet (d. 1309) November – 'Ala' al-Dawla Simnani, Persian Sufi mystic and writer (d. 1336) Albertino Mussato, Paduan statesman, poet and chronicler (d. 1329) Constantine Palaiologos, Byzantine prince and general, son of Michael VIII (d. 1306) Daniel of Moscow (Aleksandrovich), Russian prince (d. 1303) Zangpo Pal, Tibetan religious leader (d. 1323) Elizabeth of Sicily, queen consort of Hungary (House of Anjou) (d. 1303) Konoe Iemoto, Japanese nobleman (kugyō) and regent (d. 1296) Pier Saccone Tarlati di Pietramala, Italian nobleman and condottiero (d. 1356) Władysław I Łokietek ("Elbow-High"), king of Poland (d. 1333) February 28 – Henry III ("the Good"), duke of Brabant (b. 1230) April 1 – Ahi Evran, Bektashi Sufi preacher and poet (b. 1169) May 25 – Alexander IV, pope of the Catholic Church (b. 1199) July 8 – Adolf IV of Holstein, German nobleman (House of Schaumburg) July 25 – Nicephorus II of Constantinople, Byzantine patriarch August – John FitzThomas, 1st Baron Desmond, Norman Irish nobleman, killed in battle August 24 – Ela of Salisbury, English noblewoman (b. 1187) September 18 – Konrad von Hochstaden, German archbishop September 22/27 – Plaisance of Antioch, queen consort of Cyprus (b. 1235) October 27 – Sancho of Castile, Spanish archbishop (b. 1233) November 2 – Bettisia Gozzadini, Bolognese noblewoman and academic lawyer (b. 1209) November 9 – Sancha of Provence, Queen of the Romans, German queen consort (b. 1225) November 26 – Hōjō Shigetoki, Japanese samurai (b. 1198) November 27 – Athanasius III of Alexandria, Egyptian pope November 28 – Al-Mustansir II, Abbasid ruler (caliph) of Cairo, killed Abu Bakr Ibn Sayyid al-Nās, Andalusian theologian (b. 1200) An-Nasir Dawud, Kurdish ruler, Ayyubid ruler (emir) of Damascus (b. 1206) Benedict II of Esztergom, Hungarian chancellor, governor and archbishop Conrad I, Burgrave of Nuremberg ("the Pious"), German nobleman and knight (b. 1186) Qin Jiushao, Chinese mathematician and writer (b. 1202) Sayf al-Din Bakharzi, Persian poet and sheikh (b. 1190) Stephen of Bourbon, French Dominican preacher (b. 1180) ^ Steven Runciman (1952). A History of The Crusades. Vol III: The Kingdom of Acre, p. 240. ISBN 978-0-241-29877-0. ^ Bartusis, Mark C. (1997). The Late Byzantine Army: Arms and Society, 1204–1453, pp. 40–41. University of Pennsylvania Press. ISBN 0-8122-1620-2. ^ Nicol, Donald M. (1993). The Last Centuries of Byzantium, 1261–1453, p. 35 (Second ed.). Cambridge: Cambridge University Press. ISBN 978-0-521-43991-6. ^ Hackel, Sergei (2001). The Byzantine Saint, p. 71 (2001 ed.). St. Vladimir's Seminary Press. ISBN 0-88141-202-3. ^ Rossabi, Morris (1988). Khubilai Khan: His Life and Times, p. 51. Los Angeles: University of California Press. ISBN 978-0-520-06740-0. ^ Lock, Peter (2013). The Routledge Companion to the Crusades. Routledge. p. 112. ISBN 9781135131371. ^ Williams, Hywel (2005). Cassell's Chronology of World History. London: Weidenfeld & Nicolson. pp. 144–146. ISBN 0-304-35730-8. ^ BBC History, July 2011, p. 12. ^ Lawrence, C. H. (1984). "The University in State and Church". In Aston, T. H.; Catto, J. I. (eds.). The History of the University of Oxford. Vol. 1. Oxford University Press. Retrieved from " 3One hundred years, from 1101 to 1200 See also: Renaissance of the 12th century Millennia 2nd millennium Centuries 11th century 12th century 13th century Timelines 11th century 12th century 13th century State leaders 11th century 12th century 13th century Decades 1100s 1110s 1120s 1130s 1140s 1150s 1160s 1170s 1180s 1190s Categories: Births – Deaths Establishments – Disestablishments vte Eastern Hemisphere at the beginning of the 12th century The 12th century is the period from 1101 to 1200 in accordance with the Julian calendar. In the history of European culture, this period is considered part of the High Middle Ages and overlaps with what is often called the "'Golden Age' of the Cistercians". The Golden Age of Islam experienced significant development, particularly in Islamic Spain. In Song dynasty China, an invasion by Jurchens caused a political schism of north and south. The Khmer Empire of Cambodia flourished during this century, while the Fatimids of Egypt were overtaken by the Ayyubid dynasty. Following the expansions of the Ghaznavids and Ghurid Empire, the Muslim conquests in the Indian subcontinent took place at the end of the century. Main article: 1100s The Ghurid Empire converted to Islam from Buddhism. 1101: In July, the Treaty of Alton is signed between Henry I of England and his older brother Robert, Duke of Normandy in which Robert agrees to recognize Henry as king of England in exchange for a yearly stipend and other concessions. The agreement temporarily ends a crisis in the succession of the Anglo-Norman kings. 1101–1103: David the Builder takes over Kakheti and Hereti (now parts of Georgia). 1102: King Coloman unites Hungary and Croatia under the Hungarian Crown. 1102: Muslims conquer Señorío de Valencia. 1103–1104: A church council is convened by King David the Builder in Urbnisi to reorganize the Georgian Orthodox Church. 1104: In the Battle of Ertsukhi, King David the Builder defeats an army of Seljuks. 1104: King Jayawarsa of Kadiri (on Java) ascends to the throne.[citation needed] 1106: Battle of Tinchebray. 1107–1111: Sigurd I of Norway becomes the first Norwegian king to embark on a crusade to the Holy Land. He fights in Lisbon and on various Mediterranean isles and helps the King of Jerusalem to take Sidon from the Muslims. 1108: By the Treaty of Devol, signed in September, Bohemond I of Antioch has to submit to the Byzantine Empire, becoming the vassal of Alexius I. 1109: On June 10, Bertrand of Toulouse captures the County of Tripoli (northern Lebanon/western Syria). 1109: In the Battle of Nakło, Boleslaus III Wrymouth defeats the Pomeranians and re-establishes Polish access to the sea. 1109: On August 24, in the Battle of Hundsfeld, Boleslaus III Wrymouth defeats Emperor Henry V of Germany and stops German expansion eastward. Main article: 1110s 1111: On April 14, during Henry V's first expedition to Rome, he is crowned Holy Roman Emperor. 1113: Paramavishnulok is crowned as King Suryavarman II in Cambodia. He expands the Khmer Empire and builds Angkor Wat during the first half of the century. He establishes diplomatic relations with China. 1115: The Georgian army occupies Rustavi in the war with the Muslims. 1115: In Java, King Kamesvara of Kadiri ascends to the throne. Janggala ceases to exist and comes under Kadiri domination, highly possible under royal marriage. During his reign, Mpu Dharmaja writes Kakawin Smaradahana, a eulogy for the king which become the inspiration for the Panji cycle tales, which spread across Southeast Asia.[1] 1116: The Byzantine army defeats the Turks at Philomelion. 1116: Death of doña Jimena Díaz, governor of Valencia from 1099 to 1102. c. 1119: The Knights Templar are founded to protect Christian pilgrims in Jerusalem. Main article: 1120s A Black and White Photo of the 12th century Cuenca Cathedral (built from 1182 to 1270) in Cuenca, Spain 1120: On January 16, the Council of Nablus, a council of ecclesiastic and secular lords in the crusader Kingdom of Jerusalem, establishes the first written laws for the kingdom. 1120: On November 25, William Adelin, the only legitimate son of King Henry I of England, drowns in the White Ship Disaster, leading to a succession crisis which will bring down the Norman monarchy of England. 1121: On August 12, in the Battle of Didgori, the greatest military victory in Georgian history, King David the Builder with 45,000 Georgians, 15,000 Kipchak auxiliaries, 500 Alan mercenaries and 100 French Crusaders defeats a much larger Seljuk-led Muslim coalition army. 1121: On December 25, St. Norbert and 29 companions make their solemn vows in Premontre, France, establishing the Premonstratensian Order. 1122: The Battle of Beroia (Modern-day Stara Zagora, Bulgaria) results in the disappearance of the Pechenegs Turkish tribe as an independent force. 1122: On September 23, the Concordat of Worms (Pactum Calixtinum) is drawn up between Emperor Henry V and Pope Calixtus II bringing an end to the first phase of the power struggle between the papacy and the Holy Roman Empire. 1122: King David the Builder captures Tbilisi and declares it the capital city of Georgia, ending 400 years of Arab rule. 1123: The Jurchen dynasty of China forces Koryo (now Korea) to recognize their suzerainty. 1124: In April or May, David I is crowned King of the Scots. 1125: On June 11, in the Battle of Azaz, the Crusader states, led by King Baldwin II of Jerusalem, defeat the Seljuk Turks. 1125: In November, the Jurchens of the Jin dynasty declare war on the Song dynasty, beginning the Jin–Song wars. 1125: Lothair of Supplinburg, duke of Saxony, is elected Holy Roman Emperor instead of the nearest heir, Frederick of Swabia, beginning the great struggle between Guelphs and Ghibellines. 1127: The Northern Song dynasty loses power over northern China to the Jin dynasty. 1128: On June 24, the Kingdom of Portugal gains independence from the Kingdom of León at the Battle of São Mamede; (recognised by León in 1143). Main article: 1130s The temple complex of Angkor Wat, built during the reign of Suryavarman II in Cambodia of the Khmer Era. 1130–1180: 50-year drought in what is now the American Southwest. 1130–1138: Papal schism, Pope Innocent II vs. Antipope Anacletus II. 1130: On March 26, Sigurd I of Norway dies. A golden era of 95 years comes to an end for Norway as civil wars between the members of Harald Fairhair's family line rage for the remainder of the century. 1130: On Christmas Day, Roger II is crowned King of Sicily, the royal title being bestowed on him by Antipope Anacletus II. 1132: The Southern Song dynasty establishes China's first permanent standing navy, although China had a long naval history prior. The main admiral's office is at the port of Dinghai. 1132–1183: the Chinese navy increases from a mere 3,000 to 52,000 marine soldiers stationed in 20 different squadrons. During this time, hundreds of treadmill-operated paddle wheel craft are assembled for the navy to fight the Jin dynasty in the north. 1135: King Jayabaya of Kadiri ascends to the throne.[2] 1135–1154: The Anarchy takes place, during a period of civil war in England. 1136: Suger begins rebuilding the abbey church at St Denis north of Paris, which is regarded as the first major Gothic building. 1137: On July 22, the future King Louis VII of France marries Eleanor, the Duchess of Aquitaine. 1138: On October 11, the 1138 Aleppo earthquake devastates much of northern Syria. 1139: In April, the Second Lateran Council ends the papal schism. 1139: On July 5, in the Treaty of Mignano, Pope Innocent II confirms Roger II as King of Sicily, Duke of Apulia, and Prince of Capua and invests him with his titles. 1139: On July 25, the Portuguese defeat the Almoravids led by Ali ibn Yusuf in the Battle of Ourique; Prince Afonso Henriques is acclaimed King of Portugal by his soldiers. Main article: 1140s Averroes in a 14th-century painting by Andrea di Bonaiuto 1140–1150: Collapse of the Ancestral Puebloan culture at Chaco Canyon (modern-day New Mexico). 1141: The Treaty of Shaoxing ends the conflict between the Jin dynasty and Southern Song dynasty, legally establishing the boundaries of the two countries and forcing the Song dynasty to renounce all claims to its former territories north of the Huai River. The treaty reduces the Southern Song into a quasi-tributary state of the Jurchen Jin dynasty. 1143: Manuel I Komnenos is crowned as Byzantine emperor after the death of John II Komnenos. 1143: Afonso Henriques is proclaimed King of Portugal by the cortes. 1143: The Treaty of Zamora recognizes Portuguese independence from the Kingdom of León. Portugal also recognizes the suzerainty of the pope. 1144: On December 24, Edessa falls to the Atabeg Zengi. 1145–1148: The Second Crusade is launched in response to the fall of the County of Edessa. 1147: On October 25, the four-month-long Siege of Lisbon successfully brings the city under definitive Portuguese control, expelling the Moorish overlords. 1147: A new Berber dynasty, the Almohads, led by Emir Abd al-Mu'min, takes North Africa from the Almoravides and soon invades the Iberian Peninsula. The Almohads began as a religious movement to rid Islam of impurities. 1147: The Wendish Crusade against the Polabian Slavs (or "Wends") in what is now northern and eastern Germany. Main article: 1150s 1150: Ramon Berenguer IV, Count of Barcelona marries Petronilla, the Queen of Aragon. 1151: The Treaty of Tudilén is signed by Alfonso VII of León and Raymond Berengar IV, Count of Barcelona, recognizing the Aragonese conquests south of the Júcar and the right to expand in and annex the Kingdom of Murcia. 1153: The Treaty of Wallingford, ends the civil war between Empress Matilda and her cousin King Stephen of England fought over the English crown. Stephen acknowledges Matilda's son Henry of Anjou as heir. 1153: The First Treaty of Constance is signed between Emperor Frederick I and Pope Eugene III, by the terms of which, the emperor is to prevent any action by Manuel I Comnenus to reestablish the Byzantine Empire on Italian soil and to assist the pope against his enemies in revolt in Rome. 1154: the Moroccan-born Muslim geographer Muhammad al-Idrisi publishes his Geography. 1154: On December 27, Henry II is crowned King of England at Westminster Abbey. 1155: Pope Adrian IV grants overlordship of Ireland to Henry II of England in the bull Laudabiliter. 1156: On June 18, the Treaty of Benevento is entered into by Pope Adrian IV and the Norman Kingdom of Sicily. After years of turbulent relations, the popes finally settles down to peace with the Hauteville kings. The kingship of William I is recognized over all Sicily, Apulia, Calabria, Campania, and Capua. The tribute to the pope of 600 schifati agreed upon by Roger II in 1139 at Mignano is affirmed and another 400 shift is added for the new lands. 1158: The Treaty of Sahagún ends the war between Castile and León. Main article: 1160s The Liuhe Pagoda of Hangzhou, China, 1165 1161: the Song dynasty Chinese navy, employing gunpowder bombs launched from trebuchets, defeats the enormous Jin dynasty navy in the East China Sea in the Battle of Tangdao and on the Yangtze River in the Battle of Caishi. 1161: Kilij Arslan II, Sultan of Rum, makes peace with the Byzantine Empire, recognizing the emperor's primacy. 1161: In the siege of Ani, troops from the Kingdom of Georgia take control over the city, only to have it sold for the second time to the Shaddadids, a Kurdish dynasty. 1162: Genghis Khan, the founder of the Mongol Empire, is born as Temüjin in present-day Mongolia. 1163: The Norwegian Law of Succession takes effect. 1165–1182: Tensions and disputes between the Pagan Empire and the Kingdom of Polonnaruwa causes the Sinhalese under Parakramabahu the Great to raid Burma. 1168: King Valdemar I of Denmark conquers Arkona on the Island of Rügen, the strongest pagan fortress and temple in northern Europe. 1169: Political disputes within the Pandya Empire sparks the decade-long Pandyan Civil War. 1169: On May 1, the Norman invasion of Ireland begins. Richard fitzGilbert de Clare ('Strongbow') allies with the exiled Irish chief, Dermot MacMurrough, to help him recover his kingdom of Leinster. Main article: 1170s The defense of the Carroccio during the battle of Legnano (1176) by Amos Cassioli (1832–1891) 1170: The Treaty of Sahagún is signed by Alfonso VIII of Castile and Alfonso II of Aragon. Based on the terms of the accord, Alfonso VIII agrees to provide Alfonso II with three hostages, to be used as tribute payments owed by Ibn Mardanīš of Valencia and Murcia. 1170: On December 29, Thomas Becket is murdered in Canterbury Cathedral. 1171: Saladin deposes the last Fatimid Caliph Al-'Āḍid and establishes the Ayyubid dynasty. 1171: On November 11, Henry II of England lands in Ireland to assert his claim as Lord of Ireland. 1172: The Pandyan city of Madurai is sacked by the Sinhalese army due to an attempt to drive off the rival throne claimant, Kulasekara Pandyan. 1173: Sinhalese king Parakramabahu the Great gains a decisive victory by invading the Chola Empire as an ally of the Pandyas in the Pandyan Civil War. 1174: On July 12, William I of Scotland is captured by the English in the Battle of Alnwick. He accepts the feudal overlordship of the English crown and pays ceremonial allegiance at York. 1175: Hōnen Shōnin (Genkū) founds the Jōdo shū (Pure Land) sect of Buddhism. 1175: The Treaty of Windsor is signed by King Henry II of England and the High King of Ireland, Ruaidrí Ua Conchobair. 1176: On May 29, Frederick Barbarossa's forces are defeated in the Battle of Legnano by the Lombard League which results in the emperor's acknowledgment of the pope's sovereignty over the Papal States and Alexander acknowledging the emperor's overlordship of the imperial Church. 1176: On September 17, The Battle of Myriokephalon (Myriocephalum; Turkish: Miryakefalon Savaşı) is fought between the Byzantine Empire and the Seljuk Turks in Phrygia. It is a serious reversal for the Byzantine forces and will be the final, unsuccessful, effort by the Byzantines to recover the interior of Anatolia from the Seljuk Turks. 1177: The Treaty or Peace of Venice is signed by the papacy and its allies, and Frederick I, Holy Roman Emperor. The Norman Kingdom of Sicily also participates in negotiations and the treaty thereby determines the political course of all of Italy for the next several years. 1178: Chinese writer Zhou Qufei, a Guangzhou customs officer, writes of an island far west in the Indian Ocean (possibly Madagascar), from where people with skin "as black as lacquer" and with frizzy hair were captured and purchased as slaves by Arab merchants. 1179: The Treaty of Cazola (Cazorla) is signed by Alfonso II of Aragon and Alfonso VIII of Castile, dividing Andalusia into separate zones of conquest for the two kingdoms, so that the work of the Reconquista would not be stymied by internecine feuding. Main article: 1180s 1180: The Portuguese Navy defeats a Muslim fleet off the coast of Cape Espichel. 1180–1185: the Genpei War in Japan. 1181: Parakramabahu the Great conducts a large-scale raid on Burma, after a ship transporting a Sinhalese princess to the Khmer Empire is attacked by Burmese naval fleets. 1182: Religious reformations of Theravada Buddhism in Pagan Burma under the patronage of Narapatisithu are continued with the end of the Polonnaruwa-Pagan War. 1182: Revolt of the people of Constantinople against the Latins, whom they massacre, proclaiming Andronicus I Comnenus as co-emperor. 1183: On January 25, the final Peace of Constance between Frederick Barbarossa, the pope and the Lombard towns is signed, confirming the Peace of Venice of 1177. 1183: On September 24, Andronicus I Comnenus has his nephew Alexius II Comnenus strangled. 1184: On March 24, Queen Tamar, King of Georgia, accedes to the throne as sole ruler after reigning with her father, George III, for six years. 1184: Diet of Pentecost organised by Emperor Frederick I in Mainz. 1185: The Uprising of Asen and Peter against the Byzantine Empire leads to the restoration of the Bulgarian Empire. 1185: Andronicus I Comnenus is deposed and, on September 12, executed as a result of the Norman massacre of the Greeks of Thessalonika. 1185: The cathedral school (Katedralskolan) in Lund, Sweden, is founded. The school is the oldest in northern Europe and one of the oldest in all of Europe. 1185: Beginning in this year the Kamakura shogunate deprives the emperor of Japan of political power. 1186: On January 27, the future Holy Roman Emperor Henry VI marries Constance of Sicily, the heiress to the Sicilian throne. 1187: On July 4, in the Battle of Hattin, Saladin defeats the king of Jerusalem. 1187: In August, the Swedish royal and commercial center Sigtuna is attacked by raiders from Karelia, Couronia, and/or Estonia.[3] 1188: The Riah were introduced into the Habt and south of Tetouan by the Almohad caliph, Abu Yusuf Yaqub al-Mansur, and Jochem and Acem were introduced in Tamesna.[4] 1189: On September 3, Richard I is crowned King of England at Westminster. 1189: On November 11, William II of Sicily dies and is succeeded by his illegitimate cousin Tancred, Count of Lecce instead of Constance. 1189–1192: The Third Crusade is an attempt by European leaders to wrest the Holy Land from Saladin. Main articles: 1190s and 1200s Richard I of England, or Richard the Lionheart. 1190: On June 10, Emperor Frederick Barbarossa drowns in the River Salef, leaving the Crusader army under the command of the rivals Philip II of France and Richard I of England, which ultimately leads to the dissolution of the army. 1191: Holy Roman Emperor Henry VI attacked the Kingdom of Sicily from May to August but fails and withdrawn, with Empress Constance captured (released 1192). 1191: On September 7, Saladin is defeated by Richard I of England at the Battle of Arsuf. 1192: In April, Isabella I begins her reign as Christian Queen of the Kingdom of Jerusalem 1192: In the Battle of Jaffa, King Richard the Lionheart defeats Saladin. 1192: In June, the Treaty of Ramla is signed by Saladin and Richard Lionheart. Under the terms of the agreement, Jerusalem will remain under Muslim control. However, the city will be open to Christian pilgrims. The Latin Kingdom is reduced to a coastal strip that extends from Tyre to Jaffa. 1192: Minamoto no Yoritomo is appointed Sei-i Taishōgun, "barbarian-subduing great general", shōgun for short, the first military dictator to bear this title. 1192: Sultan Shahābuddin Muhammad Ghori establishes the first Muslim empire in India for 14 years (1192–1206) by defeating Prithviraj Chauhan. 1193: Nalanda, the great Indian Buddhist educational centre, is destroyed. 1194: Emperor Henry VI conquers the Kingdom of Sicily. 1195: On June 16, the struggle of Shamqori. Georgian forces annihilate the army of Abu Baqar. 1198: The brethren of the Crusader hospital in Acre are raised to a military order of knights, the Teutonic Knights, formally known as the Order of the Knights of the Hospital of St. Mary of the Teutons in Jerusalem. 1199: Pope Innocent III writes to Kaloyan, inviting him to unite the Bulgarian Church with the Roman Catholic Church. 1200: Construction begins on the Grand Village of the Natchez near Natchez, Mississippi. This ceremonial center for the Natchez people is occupied and built until the early 17th century.[5] Eastern Hemisphere at the end of the 12th century China is under the Northern Song dynasty. Early in the century, Zhang Zeduan paints Along the River During the Qingming Festival. It will later end up in the Palace Museum, Beijing. In southeast Asia, there is conflict between the Khmer Empire and the Champa. Angkor Wat is built under the Hindu king Suryavarman II. By the end of the century, the Buddhist Jayavarman VII becomes the ruler. Japan is in its Heian period. The Chōjū-jinbutsu-giga is made and attributed to Toba Sōjō. It ends up at the Kōzan-ji, Kyoto. In Oceania, the Tu'i Tonga Empire expands to a much greater area. Europe undergoes the Renaissance of the 12th century. The blast furnace for the smelting of cast iron is imported from China, appearing around Lapphyttan, Sweden, as early as 1150. Alexander Neckam is the first European to document the mariner's compass, first documented by Shen Kuo during the previous century. Christian humanism becomes a self-conscious philosophical tendency in Europe. Christianity is also introduced to Estonia, Finland, and Karelia. The first medieval universities are founded. Pierre Abelard teaches. Middle English begins to develop, and literacy begins to spread outside the Church throughout Europe.[6] In addition, churchmen are increasingly willing to take on secular roles. By the end of the century, at least a third of England's bishops also act as royal judges in secular matters.[7] The Ars antiqua period in the history of the medieval music of Western Europe begins. The earliest recorded miracle play is performed in Dunstable, England. Gothic architecture and trouvère music begin in France. During the middle of the century, the Cappella Palatina is built in Palermo, Sicily, and the Madrid Skylitzes manuscript illustrates the Synopsis of Histories by John Skylitzes. Fire and plague insurance first become available in Iceland, and the first documented outbreaks of influenza there happens. The medieval state of Serbia is formed by Stefan Nemanja and then continued by the Nemanjić dynasty. By the end of the century, both the Capetian dynasty and the House of Anjou are relying primarily on mercenaries in their militaries. Paid soldiers are available year-round, unlike knights who expected certain periods off to maintain their manor lifestyles.[8] In India, Hoysala architecture reaches its peak. In the Middle East, the icon of Theotokos of Vladimir is painted probably in Constantinople. Everything but the faces will later be retouched, and the icon will go to the Tretyakov Gallery of Moscow. The Georgian poet Shota Rustaveli composes his epic poem The Knight in the Panther's Skin. Shahab al-Din Suhrawardi founds his "school of illumination". In North Africa, the kasbah of Marrakesh is built, including the city gate Bab Agnaou and the Koutoubia mosque. In sub-Saharan Africa, Kente cloth is first woven. In France, the first piedfort coins were minted. The city of Tula burns down, marking the end of the Toltec Empire In West Africa the Ife Empire is established. See also: Timeline of historic inventions § 12th century 1104: The Venice Arsenal of Venice, Italy, is founded. It employed some 16,000 people for the mass production of sailing ships in large assembly lines, hundreds of years before the Industrial Revolution. 1106: Finished building of Gelati. 1107: The Chinese engineer Wu Deren combines the mechanical compass vehicle of the south-pointing chariot with the distance-measuring odometer device. 1111: The Chinese Donglin Academy is founded. 1165: The Liuhe Pagoda of Hangzhou, China, is built. 1170: The Roman Catholic notion of Purgatory is defined.[9] 1185: First record of windmills. Wikimedia Commons has media related to 12th century. ^ Soekmono, R, Drs., Pengantar Sejarah Kebudayaan Indonesia 2, 2nd ed. Penerbit Kanisius, Yogyakarta, 1973, 5th reprint edition in 1988 p.57 ^ Britannica, T. Editors of Encyclopaedia (1998, July 20). Kaḍiri. Encyclopedia Britannica. ^ Enn Tarvel (2007). Sigtuna hukkumine. Archived 2017-10-11 at the Wayback Machine Haridus, 2007 (7-8), p 38–41 ^ Notice sur les Arabes hilaliens. Ismaël Hamet. p. 248. ^ Francine Weiss and Mark R. Barnes (May 3, 1989). "National Register of Historic Places Registration: Grand Village of the Natchez Site / Fatherland Plantation Site (22-Ad-501)" (pdf). National Park Service. and Accompanying 3 photos, from 1989. (680 KB) ^ Warren 1961, p. 129. ^ Warren 1961, p. 159. ^ Warren 1961, p. 60-61. ^ Le Goff, Jacques (1986). The Birth of Purgatory. Chicago: University of Chicago Press. ISBN 0226470822. Warren, Wilfred Lewis (1961). King John. University of California Press. p. 362. ISBN 9780520036437. {{cite book}}: ISBN / Date incompatibility (help) Retrieved from " 4 The following pages link to 12th century External tools (link count transclusion count sorted list) · See help page for transcluding these entries Showing 50 items. View (previous 50 | next 50) (20 | 50 | 100 | 250 | 500)Antisemitism in Christianity (links | edit) Catharism (links | edit) List of decades, centuries, and millennia (links | edit) Dialect (links | edit) House of Habsburg (links | edit) House of Hohenzollern (links | edit) History of Mali (links | edit) Maimonides (links | edit) Passport (links | edit) Sumba (links | edit) Taoism (links | edit) Weregild (links | edit) Zanzibar (links | edit) 20th century (links | edit) 1040 (links | edit) 15th century (links | edit) 16th century (links | edit) 17th century (links | edit) 18th century (links | edit) 14th century (links | edit) 1st century (links | edit) 13th century (links | edit) 4th century (links | edit) 11th century (links | edit) 1282 (links | edit) 7th century (links | edit) 10th century (links | edit) 9th century (links | edit) 8th century (links | edit) 6th century (links | edit) 5th century (links | edit) 3rd century (links | edit) 2nd century (links | edit) 4th century BC (links | edit) 1st century BC (links | edit) 2nd century BC (links | edit) 3rd century BC (links | edit) 5th century BC (links | edit) 6th century BC (links | edit) 21st century BC (links | edit) 11th century BC (links | edit) 1000s (decade) (links | edit) 1040s (links | edit) 1299 (links | edit) 1154 (links | edit) 1160s (links | edit) 1163 (links | edit) 1141 (links | edit) 1135 (links | edit) 1204 (links | edit) View (previous 50 | next 50) (20 | 50 | 100 | 250 | 500) Retrieved from " WhatLinksHere/12th_century" Generalized linear models (GLMs) stand as a cornerstone in the field of statistical analysis, extending the concepts of traditional linear regression to accommodate various types of response variables. These models allow researchers and data analysts to accurately describe and infer the relationships between a dependent variable and one or more independent variables, especially when the response variable does not adhere to normal distribution assumptions. This flexibility makes GLMs invaluable across a multitude of disciplines, including economics, medicine, engineering, and the social sciences. Throughout this article, we will delve into the components, types, and applications of GLMs, offering insights into their theoretical underpinnings and practical uses. Whether you are a student stepping into the world of statistics, a data analyst looking to refine your modeling techniques, or a researcher conducting complex analyses, understanding GLMs will enhance your ability to extract meaningful insights from data. Historical Background and Development of GLMs The journey of generalized linear models (GLMs) began in the 20th century, marking a significant milestone in the evolution of statistical methodologies. The concept of GLMs was formally introduced in 1972 in a groundbreaking paper that extended the classical linear regression model. Their work was revolutionary, providing a comprehensive framework that unified various statistical models under a single theory. This framework was designed to accommodate response variables with error distribution models other than the normal distribution, thus broadening the applicability of regression techniques significantly. Before the advent of GLMs, statistical analysis was somewhat fragmented, with different models and methods developed independently for specific types of data and analysis objectives. Classical linear regression models were well-suited for continuous data following a normal distribution. However, they fell short when dealing with binary, count, or other types of non-normally distributed data. The development of GLMs addressed these limitations by introducing flexibility in both the distribution of the response variable and the link between the response and the predictors. Comparison with Classical Linear Regression Models: Classical linear regression models assume that the response variable is normally distributed (or that the errors in the prediction are). GLMs allow for the response variable to follow various distributions from the exponential family, such as binomial, Poisson, and gamma, accommodating a wider range of data types. Classical linear models rely on a linear relationship between the dependent and independent variables. GLMs introduce a link function, enabling the modeling of non-linear relationships within a linear framework, thereby offering greater flexibility in analyzing the relationship between variables. Classical linear regression models are limited in their application to datasets that strictly meet the assumptions of linearity, independence, and normality of residuals. GLMs provide a more versatile framework that can be tailored to suit various types of data and analysis needs, overcoming the limitations of classical linear regression. Classical Linear Regression Models are best suited for continuous data that fits the normal distribution. GLMs are designed to handle a broader scope of data types, including binary, count, and continuous data, making them applicable in a wider array of disciplines and research scenarios. Classical Linear Regression Models assume homoscedasticity (constant variance of error terms) across observations. GLMs allow for different error structures, accommodating heteroscedasticity and other complexities in the data and providing a more accurate representation of the variance. GLMs are based on a mathematical framework that assumes the response variable's distribution comes from the exponential family. This family includes a wide range of distributions, such as normal, binomial, Poisson, and gamma, among others. The choice of distribution depends on the nature of the data being modeled (e.g., continuous, binary, count data), which allows GLMs to be tailored to specific types of data. A probability distribution belongs to the exponential family if its probability density function (pdf) or probability mass function (pmf) can be expressed in the following form: Core components of GLMs: At the heart of GLMs are three core components that define their structure and functionality: the error distribution, the linear predictor, and the link function. Understanding these components is crucial for effectively applying GLMs to various data types and research questions. 1. Error Distribution (Family) The error distribution, often referred to as the "family," specifies the probability distribution of the response variable. This component is foundational in GLMs because it determines the form of the likelihood function used for parameter estimation. The choice of distribution is guided by the nature of the response variable (e.g., continuous, binary, count) and the model's appropriateness and accuracy. The exponential family of distributions provides a flexible set of options for GLMs, including but not limited to: Normal Distribution: Used for continuous data where residuals are assumed to be normally distributed. Binomial Distribution: Suitable for binary response variables, such as in logistic regression, where the outcome is success or failure, yes or no. Poisson Distribution: Employed for count data where the response variable represents the number of events occurring in a fixed interval or space, as in Poisson regression. Gamma Distribution: Used for positive continuous data, often applied in models where the response variable represents time until an event occurs. 2. Linear Predictor The linear predictor is the component of the GLM that incorporates the independent variables (predictors). It is a linear combination of these variables, each multiplied by a coefficient that quantifies the relationship between the predictor and the response variable. The linear predictor is expressed as: 3. Link Function The link function is a critical component that connects the linear predictor to the mean (μ) of the response variable's distribution. It transforms the expected value of the response variable to the scale of the linear predictor, ensuring that the model's predictions stay within the range that is appropriate for the response variable's distribution. The choice of link function depends on the error distribution and can significantly influence the interpretation and performance of the model. Common link functions include: Types of GLMs: Distribution: Binomial Link Function: Logit Usage: Logistic regression is used for binary outcomes (e.g., success/failure, yes/no). It models the probability that an observation falls into one of the two categories. Logistic regression is widely applied in fields like medicine for disease diagnosis, marketing for customer churn prediction, and finance for credit risk assessment. 2. Poisson Regression Distribution: Poisson Link Function: Log Usage: Poisson regression is used for count data, particularly when modeling the number of times an event occurs in a fixed period of time or space. It is commonly used in public health for incidence rate analysis, in insurance for claim frequency modeling, and in traffic engineering for accident count prediction. 3. Multinomial Logistic Regression Distribution: Multinomial Link Function: Softmax 4. Ordinal Regression Distribution: Ordinal Link Function: Various (e.g., logit, probit) Usage: Ordinal regression models outcomes with a natural order but without a consistent difference between adjacent categories. This type is used in survey analysis (e.g., rating scales), education (e.g., grading systems), and any scenario where responses fall into naturally ordered categories. 5. Negative Binomial Regression Distribution: Negative Binomial Link Function: Log Usage: Negative binomial regression is an extension of Poisson regression, used when the data exhibit overdispersion (i.e., the variance is greater than the mean). It's useful in ecological data for species counts, in public health for disease count data, and anywhere the assumption of equal mean and variance is violated. 6. Gamma Regression Distribution: Gamma Link Function: Inverse Usage: Gamma regression models positive continuous data, particularly for skewed distributions. It's applicable in modeling time-to-event data, such as survival times in medical research, or service times in operations research. 7. Zero-Inflated Models Distribution: Poisson, Negative Binomial, or others with zero-inflation component Link Function: Log (for count part) Usage: Zero-inflated models are used when the data contain excess zeros beyond what the standard distribution would predict. They're useful in insurance for modeling claim counts with many policyholders having zero claims, in environmental science for species abundance data, and in healthcare for count data on rare events. Model Estimation and Fitting in Generalized Linear Models (GLMs) Maximum Likelihood Estimation is a method used to estimate the parameters of a statistical model. In the context of GLMs, MLE seeks to find the parameter values that maximize the likelihood function, which represents the probability of observing the given data under the assumed model. The likelihood function depends on the chosen distribution of the response variable. Equation: The general form of the likelihood equation for a GLM can be complex and varies with the distribution. However, the goal remains to maximize L(β|data), the probability of the data given the parameters. Process: Maximizing the log-likelihood is often more practical due to the additive nature of the logarithm, which simplifies the calculation, especially with computational tools. The process involves taking the derivative of the log-likelihood function with respect to each parameter in θ, setting these derivatives (partial derivatives) equal to zero, and solving for the parameters. This results in the maximum likelihood estimates (MLEs) of the parameters. However, for many GLMs, these equations do not have closed-form solutions and must be solved iteratively. Iteratively reweighted least squares (IRLS) The goal of IRLS is to minimize the difference between the observed responses and the model's predictions, while accounting for the non-linear link function and the variance structure of the distribution. The process involves several key steps, iteratively repeated until convergence: Start with an Initial Estimate: Begin with an initial guess for the parameters. This can be as simple as the coefficients from a linear regression if applicable or zeros. Compute the Weights: Based on the current estimate of the parameters, calculate weights for each observation. These weights account for the variance of each observation's contribution to the overall model, with the variance function being a characteristic of the exponential family distribution used. Solve a Weighted Least Squares Problem: Use the current weights and the observed data to solve a weighted least squares problem. This involves computing a new estimate of the parameters that minimizes the weighted sum of squared differences between the observed responses and the predictions made by the linear predictor. Update Estimates and Repeat: The solution from the weighted least squares problem provides a new estimate of the parameters. This new estimate is then used to repeat the process, recalculating weights and solving a new weighted least squares problem, until the parameter estimates converge within a specified tolerance level. Mathematical Formulation In the context of GLMs, let's denote the linear predictor as η=Xβ, where X is the matrix of predictors, and β represents the coefficients. The iterative step in IRLS can be mathematically described by updating the estimate of β as follows: where W is a diagonal matrix of weights, which depend on the current estimate of β, and reflecting the variance of each observation's response. z is the working response variable, a transformation of the observed response that incorporates the current residuals and the link function. Evaluating and selecting the appropriate Model: It is a critical step in statistical modeling process, ensuring that the chosen model adequately represents the data and supports reliable inference. The process involves assessing the model's fit, understanding its predictive performance, and comparing it against alternative models. Here are some key concepts and techniques used in the evaluation and selection of GLMs: 1. Goodness of Fit Deviance: A measure of the goodness of fit of a GLM, deviance is the difference between the log-likelihood of the model and the log-likelihood of a saturated model (a model with a perfect fit). Lower deviance indicates a better fit, but it's crucial to compare deviance relative to other models for the same data. Pearson Chi-Squared Statistic: This metric compares the discrepancy between observed and expected values, scaled by the variance. It's another tool for assessing the fit, especially useful in count data models like Poisson regression. 2. Model Comparison and Selection Criteria Akaike Information Criterion (AIC): AIC is a widely used criterion for model selection among a set of candidate models. It balances the goodness of fit with the complexity of the model, penalizing models with more parameters to prevent overfitting. The model with the lowest AIC is typically preferred. Bayesian Information Criterion (BIC): Similar to AIC, the BIC also provides a balance between fit and complexity but imposes a stricter penalty for the number of parameters. It's particularly useful in model selection with large datasets. 3. Residual Analysis Analyzing residuals, the differences between observed and predicted values, can reveal patterns that indicate model misspecification, such as non-linearity or heteroscedasticity. Plotting residuals against fitted values or predictors can help diagnose issues with the model fit. 4. Cross-Validation Cross-validation, particularly k-fold cross-validation, is a robust method for assessing a model's predictive performance. By dividing the data into k subsets and using each in turn as a validation set while training the model on the remaining data, cross-validation provides a reliable estimate of the model's predictive accuracy on unseen data. 5. Overdispersion Particularly relevant in count models like Poisson regression, overdispersion occurs when the observed variance exceeds the mean. Ignoring overdispersion can lead to underestimated standard errors and overconfident inferences. Negative binomial regression or quasi-likelihood models are alternatives that can address overdispersion. 6. Coefficient Significance Assessing the significance of the model coefficients through Wald tests or likelihood ratio tests provides insight into the relevance of each predictor. Coefficients with p-values below a certain threshold (commonly 0.05) are typically considered statistically significant. 7. Predictive Performance Metrics For models aimed at prediction, such as logistic regression, metrics like the area under the Receiver Operating Characteristic (ROC) curve (AUC) for binary outcomes, or mean squared error (MSE) for continuous outcomes, offer quantitative measures of predictive accuracy. Common metrics for classification tasks: Common metrics for regression Implementing Generalized Linear Models (GLMs) is facilitated by various software tools and programming languages that offer built-in functions and libraries specifically designed for statistical analysis. Here, we focus on two widely used programming environments for statistical computing and machine learning: R and Python. Each environment provides powerful packages and functions for fitting, analyzing, and visualizing GLMs. R and the glm() Function R, a programming language and environment specifically designed for statistical computing and graphics, offers a straightforward and efficient way to fit GLMs through the glm() function available in the base package. This function supports various distributions and link functions, making it adaptable to different types of GLMs. Example: Logistic Regression in R # Load necessary library library(MASS) # For the dataset # Load the Pima Indians Diabetes dataset data(Pima.tr, package = "MASS") # Fit a logistic regression model to predict diabetes status # using pregnancy count, plasma glucose concentration, and BMI glm_model 0, else 0 data['affair'] = (data.affairs > 0).astype(int) # Define the model model = smf.glm('affair ~ rate_marriage + age + yrs_married + children + religious + educ', data=data, family=sm.families.Binomial()) # Fit the model result = model.fit() # Print the summary result_summary = result.summary() result_summary output: Challenges and Considerations in Using GLMs Using Generalized Linear Models (GLMs) can significantly enhance the analysis of various data types across multiple disciplines. However, as with any statistical modeling technique, there are challenges and considerations that analysts must navigate to ensure accurate, reliable results. Here are some key challenges and considerations when using GLMs: 1. Choice of Link Function and Distribution Critical Decision: The choice of link function and distribution is crucial as it directly affects the model's interpretation and performance. An inappropriate choice can lead to poor model fit, biased estimates, and incorrect inferences. Consideration: The selection should be based on the nature of the dependent variable (e.g., continuous, binary, count) and the relationship between the dependent and independent variables. 2. Model Overdispersion Issue: Overdispersion occurs when the observed variance is greater than what the model expects under the assumed distribution (e.g., Poisson distribution in count data).Impact: It can lead to underestimated standard errors and, consequently, inflated significance levels for predictors.Solution: Consider alternative distributions (e.g., Negative Binomial for count data) or use quasi-likelihood methods.3. Multicollinearity Among PredictorsProblem: High correlation among independent variables can destabilize the model, leading to inflated standard errors and unreliable coefficient estimates.Detection: Use variance inflation factors (VIFs) or correlation matrices to detect multicollinearity.Mitigation: Simplify the model by removing highly correlated predictors, or use dimensionality reduction techniques.4. Non-linearityChallenge: The linear predictor in GLMs assumes a linear relationship between the transformed mean of the response and the predictors. However, real-world data might exhibit non-linear relationships.Solution: Include polynomial terms or interaction terms in the model, or consider using Generalized Additive Models (GAMs) for more flexibility in capturing non-linear trends.5. Zero-Inflation in Count DataObservation: Excessive zeros in count data, more than what the Poisson or Negative Binomial distributions predict.Approach: Use zero-inflated models or hurdle models that explicitly model the excess zeros as part of the process.6. Missing Data and OutliersImpact: Missing values and outliers can significantly affect model estimates and predictions.Handling Strategies: Employ imputation techniques for missing data, and investigate outliers to decide whether to include, exclude, or transform them.7. Interpretability and ComplexityBalance: While GLMs can be extended to accommodate complex relationships through transformations and link functions, this may come at the cost of interpretability.Consideration: Ensure that the complexity of the model is warranted by the data structure and analysis objectives, and that the model can be reasonably interpreted and explained.8. Computational ChallengesConcern: Fitting GLMs, especially with large datasets or complex link functions, can be computationally intensive.Solution: Utilize efficient software and algorithms optimized for GLMs, and consider simplifying the model if computational limitations are a concern.ConclusionAs data continue to grow in complexity and volume, the role of GLMs in statistical analysis and predictive modeling will undoubtedly remain significant. The continuous development of statistical software and computational tools further expands the accessibility and application of GLMs across various disciplines. Despite the challenges and considerations inherent in their use, the careful application of GLMs can lead to insightful, accurate, and meaningful analyses. Generalized linear models (GLMs) are a powerful tool for data scientists, providing a flexible way to model data. In this post, you will learn about the concepts of generalized linear models (GLM) with the help of Python examples. It is very important for data scientists to understand the concepts of generalized linear models and how are they different from general linear models such as regression or ANOVA models. Generalized linear models (GLM) are a type of statistical models that can be used to model data that is not normally distributed. It is a flexible general framework that can be used to build many types of regression models, including linear regression, logistic regression, and Poisson regression. GLM provides a way to model dependent variables that have non-normal distributions. GLM also allows for the incorporation of predictor variables that are not Normally distributed. GLMs are similar to linear regression models, but they can be used with data that has a non-normal distribution. This makes GLMs a more versatile tool than linear regression models. This is different from the general linear models (linear regression / ANOVA) where response variable, Y, and the random error term ([latex]\epsilon[/latex]) have to be based solely on the normal distribution. Linear models can be expressed in terms of expected value (mean) of response variable as the following: [latex]\Large g(\mu)= \sum\limits_{i=1}^n \beta_iX_i[/latex] ... where [latex]\mu[/latex] can be expressed as E(Y) aka expected value of response variable Y. . Based on the probability distribution of the response variable, different link functions get used which transforms [latex]g(\mu)[/latex] appropriately to the output value which gets modeled using different types of regression models. If the response variable is normally distributed, the link function is identity function and the model looks like the following. Y, in the equation below, represents the expected value or E(Y). [latex]\Large Y = \sum\limits_{i=1}^n \beta_iX_i[/latex] GLM can be used for both regression and classification problems. In regression, the goal is to predict the value of the dependent variable (e.g., price of a house). In classification, the goal is to predict the class label of the dependent variable (e.g., whether a patient has cancer). Before getting into generalized linear models, lets quickly understand the concepts of general linear models. General linear models are the models which is used to predict the value of continuous response variable and, at any given predictor value, X, the response variable, Y, and random error term ([latex]\epsilon[/latex]) follows a normal distribution. The parameter of such normal distribution represents the mean as linear combination of weights (W) and predictor variable (X), and, the standard deviation of [latex]\sigma[/latex]. Linear regression and ANOVA models represent the general linear models. The diagram given below represents the same in form of simple linear regression model where there is just one coefficient. Fig 1. Response variable and error term follows normal distribution For linear regression models, the link function is identity function. Recall that a link function transforms the probabilities of the levels of a categorical response variable to a continuous scale that is unbounded. Once the transformation is complete, the relationship between the predictors and the response can be modeled with linear regression. Thus, the linear combination of weights and predictor variable is modelled as output. The linear regression models using identity function as link function can be understood as the following: [latex]\Large Y_{actual} = Y_{predicted} + \epsilon[/latex] ... The actual value is the sum of predicted value and the random error term (which can be on either side of the mean value – response value) . [latex]\Large Y_{actual} = \sum\limits_{i=1}^n \beta_iX_i + \epsilon[/latex] . As part of training regression models, one must understand that what is actually modelled is the mean of the response value and not the actual values. As the response variable Y follows normal distribution, the summation of weights and predictor variable can be equated as the expected value of Y. [latex]\Large E(Y) = \sum\limits_{i=1}^n \beta_iX_i[/latex] . The actual value of Y can be represented as the following in terms of outcome from regression model and the random error term: [latex]\Large Y_{actual} = E(Y) + \epsilon[/latex] . For the linear regression model, the identity function is link function used to model the expected value of response variable, Y, and the summation of weights and predictor variable. Thus, the g(E(Y)) becomes E(Y) which is represented as [latex]Y_{predicted}[/latex]. Thus, linear regression model (also, at times termed as general linear models) is represented as the following: [latex]\Large Y_{predicted} = \sum\limits_{i=1}^n \beta_iX_i[/latex] Given above, lets understand what are generalized linear models. In generalized linear models, the link function used to model the response variable as a function of the predictor variables are the following. Note that the Y represents the mean or expected value of the response variable. Log Link: Log(Y) – Models the logarithm of mean Y. Thus, the regression model can be represented as the following: [latex]\Large Log(Y) = \sum\limits_{i=1}^n \beta_iX_i[/latex] ... This type of regression model is called as Poisson regression. It is used to model the non-negative count value as in Poisson probability distribution. . Logit Link: [latex]\Large Log(\frac{Y}{(1 - Y)})[/latex] – Models the logarithm of odds of the probability of binary outcome. Thus, the regression model can be represented as the following: [latex]\Large Log(\frac{Y}{(1 - Y)})= \sum\limits_{i=1}^n \beta_iX_i[/latex] .. This type of regression model is called as Logistic regression. It is used to model the probability of the binary outcome. . There are other link functions which result in modeling response variable based on gamma distribution (GammaRegressor), Tweedie distribution (TweedieRegressor) etc. The above regression models used for modeling response variable with Poisson, Gamma, Tweedie distribution etc are called as Generalized Linear Models (GLM). Here are some real-world examples where generalized linear models can be used to predict continuous response variables based on their probability distribution. The table consists of reference to the SKlearn class which can be used to model the response variables. Problem TypeExampleRegression Model (Sklearn)Agriculture / weather modelingNumber of rain events per yearPoissonRegressorAgriculture / weather modelingAmount of rainfall per rainfall eventGammaRegressorAgriculture / weather modelingTotal rainfall per yearTweedieRegressorRisk modeling / insurance policy pricingNo of claim events / policyholder per yearPoissonRegressorRisk modeling / insurance policy pricingCost per claim eventGammaRegressorRisk modeling / insurance policy pricingTotal cost per policyholderTweedieRegressorPredictive maintenanceNumber of production interruption events per year PoissonRegressorPredictive maintenanceDuration of interruptionGammaRegressorPredictive maintenanceTotal interruption time per yearTweedieRegressor Generalized linear models of different kinds are used based on the probability distribution of the response variables. GLM is particularly useful when the response variable is not normally distributed or when the relationship between the predictor variables and the response variable is non-linear.If the response variable values are counts (non-negative integer valued) or relative frequencies (non-negative), Poisson regression with log-link is used. Sklearn PoissonRegressor can be used to model such response variable.If the response variable values are positive valued and skewed, Sklearn GammaRegressor with log-link can be tried.If the response variable values seem to be heavier tailed than a Gamma distribution, one may try an Inverse Gaussian distribution based regressor such as Sklearn TweedieRegressor.If the response variable is related to binary outcome, Logistic regression with Logit link can be used. Sklearn LogisticRegression can be used to model the response variable having binary outcome. GLM allows for the inclusion of interaction terms and polynomial terms. As a result, GLM is a versatile tool that can be used in a wide range of situations. Here is the summary of what you learned in this post in relation to generalized linear models: General linear models are used to model the response variable as a linear combination of weights and predictor variable given the response variable and the error term follows the exponential family of distributions.GLM can model response variable which follows distribution such as normal, Poisson, Gamma, Tweedie, binomial etc. Python Sklearn provides classes to train GLM models depending upon the probability distribution followed by the response variable. I have been recently working in the area of Data analytics including Data Science and Machine Learning / Deep Learning. I am also passionate about different technologies including programming languages such as Java/JEE, Javascript, Python, R, Julia, etc, and technologies such as Blockchain, mobile computing, cloud-native technologies, application security, cloud computing platforms, big data, etc. I would love to connect with you on Linkedin. Check out my latest book titled as First Principles Thinking: Building winning products using first principles thinking. This chapter will help you to: Explain the goals of the general linear model Describe the general linear model as an equation for a line List which tests of statistical inference fit in the general linear model The general linear model (or GLM for short) is a popular approach to statistical inference because of its versatility. The term "general" suggests that it can be applied in most cases. We'll discuss which cases are appropriate for the general linear model, how it implements the tests of significance, and the goals of the approach. In this course, we will use the GLM procedure in SPSS for each test of significance. This will limit the number of procedures and menus to learn and facilitate the inclusion of chart creation and assumption checking. The main goal of GLM is to form population estimates of the strength and direction of relationships among predictor variables and outcome variables. The estimates are usually in the form of change. That is, how much does the outcome variable change when the values of the predictor variables change. Of course, anytime we form an estimate for the population, we must account for sampling error. We want to know if our estimate is reliable (i.e., likely to be found in other samples of similar size) or just a fluke. We can form confidence intervals and perform null hypothesis significance testing on these estimates to do help with this. In addition to testing each estimate, we can assess the fit of the full model. We can summarize the extent to which our model correctly predicts the variability in the outcome variables. Although beyond the scope of this course, this feature of GLM allows us to compare different models so that we can choose the best fitting model before moving to interpretation and application of the model. The general linear model is a way to state the direction and strength of linear relationships among variables. It is a model because it is a guess about how the population values are related that is built from sample data. Just as an engineer might construct a small scale model to test hypotheses, so to does a statistician construct a model to test hypotheses about a larger population. There are many ways to state predicted relationships among variables but perhaps the most popular model is the linear model, which follows a specific form. \[Y_i = B X_i+\epsilon_i \] This is the equation for a line where \(Y_i\) is the outcome variable value for the \(i^{\textrm{th}}\) participant, B is the slope of the line, \(X_i\) is the predictor variable for the \(i^{\textrm{th}}\) participant, and \(\epsilon_i\) is the error term (more on this shortly). Slope is the change in the y-variable over the change in the x-variable. It represents both the direction (positive or negative) and strength (amount of change) of the relationship of the two variables. The slope is kind of like the correlation coefficient except that the correlation coefficient is standardized and this one does not have units. If you are using the general linear model for regression, you can easily plug in a value of x to get a value of y. If you are interested in testing statistically significant change from one group to another, a standardized slope may be easier to work with. A standardized slope will be shown as a beta (\(\beta\)). The error term is the difference between what the model predicts (e.g., multiplying \(X_i\) by the slope) and the actual outcome variable value. In null hypothesis significance testing, we assume that the slope is 0 (i.e., no change in outcome variable values across values of the predictor variable). We'll highlight the similarity of this approach with t-tests in the next section. Although the linear regression model pre-dates the general linear model, it turns out to be just one specific case or implementation of the general linear model. With the more general form, we can include multiple predictors and multiple outcomes. The formula looks very similar but the letters represent something different. \[Y = B X + U \] Here's what changed. Rather than \(Y_i\), which represented the observations of a single outcome variable, \(Y\) is a matrix of observations for multiple outcome variables. Similarly, \(X_i\) is replaced with the more general matrix of predictor variable values, \(X\). \(U\) is the matrix of errors. A Matrix is a way to organize numbers into columns and rows where the columns typically represent some grouping factor. In statistics, the columns often represent variables. By changing the number of and type of variables that give rise to \(Y\) and \(X\), we can derive the various models that we'll cover in this course. Each of the following can be modeled in the general linear model Independent samples t-test Paired Samples t-test One-way Between-Subjects ANOVA Factorial Between-Subjects ANOVA One-Way Within-Subjects ANOVA Factorial Within-Subjects ANOVA Mixed Factorial ANOVA We will not need to worry about how to change the general formula to fit analysis in this course because we will follow similar steps each time and SPSS will set up the equations for us. However, if you are particularly curious about the magic behind the scenes, this site gives a nice overview of a few of the more simple tests. It also relates these to "non-parametric" versions of these tests (in case you do not have normally distributed data). SPSS provides many ways to reach the same end for analyses. Take an independent samples t-test for example. You could run the "independent samples t-test" analysis from the "Compare Means" menu. You could dummy code your data and enter into the "Linear" analysis from the "Regression" menu. As we'll focus on this semester, you can also use the "Univariate" analysis under the "General Linear Model" menu. You will get different output from each of these but the interpretation / conclusion will be the same in each case. I prefer to use and teach GLM in SPSS because it is nicely bundles together assumption checking, figure plotting, and tests of significance / confidence intervals into one set of steps. Furthermore, once you learn the basics of how to use GLM in SPSS, you can fit most data obtained through experiments. To be clear, GLM cannot handle all data. We'll need to meet for a different course to discuss the generalized linear model and mixed effects modeling. For the rest of the chapters, I'll provide you with a step-by-step guide on how to use the GLM in each circumstance. I'm sure that, by the end, it will seem a bit redundant. That's a good thing! In this blog, you will get to know the essential mathematical topics you need to cover to become good at AI & machine learning. These topics are grouped under four core areas including linear algebra, calculus, multivariate calculus and probability theory & statistics. Linear Algebra Linear algebra is arguably the most important mathematical foundation for machine learning. At its core, machine learning is about manipulating large datasets, and linear algebra provides the tools to do this efficiently. Vector Spaces and Operations Matrices: Your Data's Best Friend Eigenvalues and Eigenvectors Matrix Decompositions Calculus Machine learning is fundamentally about optimization – finding the best parameters that minimize error in the loss function. … Continue reading This blog represents a list of questions you can ask when thinking like a product leader. The list of questions are designed to help you build your product thinking skills. … Continue reading This blog represents a list of questions you can ask when thinking like a product leader. This is a list of questions that business leaders should consider regarding business problem & opportunity, getting business impact clarity, understanding different aspects of user experience & design, success & measurement criteria, etc. If you are starting on with an exercise to solve a business problem and therefore exploring and selecting a product, you might find these questions useful. Core Problem & Opportunity Product leaders must first establish a crystal-clear understanding of the problem they're solving and the opportunity it represents. This foundational work involves deeply understanding your target users, their pain points, and the real-world scenarios where your solution will add value. … Continue reading AI agents are autonomous systems combining three core components: a reasoning engine (powered by LLM), tools for external actions, and memory to maintain context. Unlike traditional AI-powered chatbots (created using DialogFlow, AWS Lex), agents can interact with end user based on planning multi-step workflows, use specialized tools, and make decisions based on previous results. In this blog, we will learn about different agentic AI design patterns. The blog represents Python code examples to explain each of the approaches for creating AI agents. Before getting into the blog, lets quickly look at the set up code which will be used for code used in the approaches. To explain different approaches, … Continue reading Artificial Intelligence (AI) has evolved significantly, from its early days of symbolic reasoning to the emergence of large language models that rely on internet-scale data. Now, a new frontier is taking shape— Embodied AI that leverages Agentic AI. These systems move beyond static data processing to actively interact with and learn from the real world. Embodied AI, in particular, refers to intelligent agentic AI systems with physical presence—robots, drones, humanoids—that sense, reason, and act in physical environments. Together with Agentic AI, which emphasizes autonomy, goal-directed behavior, and decision-making over time, these developments represent a shift toward more dynamic, adaptive, and human-like forms of intelligence that integrate perception, cognition, and action. … Continue reading Posted in agentic ai. Tagged with agentic ai. Last updated: 25th Jan, 2025 Have you ever wondered how to seamlessly integrate the vast knowledge of Large Language Models (LLMs) with the specificity of domain-specific knowledge stored in file storage, relational databases, graph databases, vector databases, etc? As the world of LLMs continues to evolve, the need for more sophisticated and contextually relevant responses from LLMs becomes paramount. Lack of contextual knowledge can result in LLM hallucination thereby producing inaccurate, unsafe, and factually incorrect responses. This is where context engineering and prompts is used for contextually sensitive answer generation with LLMs, and, the retrieval-augmented generation method, comes into the picture. For data scientists and product managers keen … Continue reading Hey there! As I venture into building agentic MEAN apps with LangChain.js, I wanted to take a step back and revisit the core concepts of the MEAN stack. LangChain.js brings AI-powered automation and reasoning capabilities, enabling the development of agentic AI applications such as intelligent chatbots, automated customer support systems, AI-driven recommendation engines, and data analysis pipelines. Understanding how it integrates into the MEAN stack is essential for leveraging its full potential in creating these advanced AI applications. So, I put together this quick learning blog to share what I've revisited. The MEAN stack is a popular full-stack JavaScript framework that consists of MongoDB, Express.js, Angular, and Node.js. Each component plays … Continue reading Software-as-a-Service (SaaS) providers have long relied on traditional chatbot solutions like AWS Lex and Google Dialogflow to automate customer interactions. These platforms required extensive configuration of intents, utterances, and dialog flows, which made building and maintaining chatbots complex and time-consuming. The need for manual intent classification and rule-based conversation logic often resulted in rigid and limited chatbot experiences, unable to handle dynamic user queries effectively. With the advent of generative AI, SaaS providers are increasingly adopting Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), and multi-agent frameworks such as LangChain, LangGraph, and LangSmith to create more scalable and intelligent AI-driven chatbots. This blog explores how SaaS providers can leverage these technologies … Continue reading Retrieval-Augmented Generation (RAG) is an innovative generative AI method that combines retrieval-based search with large language models (LLMs) to enhance response accuracy and contextual relevance. Unlike traditional retrieval systems that return existing documents or generative models that rely solely on pre-trained knowledge, RAG technique dynamically integrates context as retrieved information related to query with LLM outputs. LangGraph, an advanced extension of LangChain, provides a structured workflow for developing RAG applications. This guide will walk through the process of building a RAG system using LangGraph with example implementations. Setting Up the Environment To get started, we need to install the necessary dependencies. The following commands will ensure that all required LangChain … Continue reading The combination of Retrieval-Augmented Generation (RAG) and powerful language models enables the development of sophisticated applications that leverage large datasets to answer questions effectively. In this blog, we will explore the steps to build an LLM RAG application using LangChain. Prerequisites Before diving into the implementation, ensure you have the required libraries installed. Execute the following command to install the necessary packages: Setting Up Environment Variables LangChain integrates with various APIs to enable tracing and embedding generation, which are crucial for debugging workflows and creating compact numerical representations of text data for efficient retrieval and processing in RAG applications. Set up the required environment variables for LangChain and OpenAI: Step … Continue reading Have you ever wondered how to use OpenAI APIs to create custom chatbots? With advancements in large language models (LLMs), anyone can develop intelligent, customized chatbots tailored to specific needs. In this blog, we'll explore how LangChain and OpenAI LLMs work together to help you build your own AI-driven chatbot from scratch. Prerequisites Before getting started, ensure you have Python (version 3.8 or later) installed and the required libraries. You can install the necessary packages using the following command: Setting Up OpenAI API Key To use OpenAI's services, you need an API key, which you can obtain by signing up at OpenAI's website (OpenAI) and generating a key from the … Continue reading When building a Retrieval-Augmented Generation (RAG) application powered by Large Language Models (LLMs), which combine the ability to generate human-like text with advanced retrieval mechanisms for precise and contextually relevant information, effective indexing plays a pivotal role. It ensures that only the most contextually relevant data is retrieved and fed into the LLM, improving the quality and accuracy of the generated responses. This process reduces noise, optimizes token usage, and directly impacts the application's ability to handle large datasets efficiently. RAG applications combine the generative capabilities of LLMs with information retrieval, making them ideal for tasks such as question-answering, summarization, or domain-specific problem-solving. This blog will walk you through the … Continue reading Artificial Intelligence (AI) agents have started becoming an integral part of our lives. Imagine asking your virtual assistant whether you need an umbrella tomorrow, or having it remind you of an important meeting—these agents now help us with weather forecasts, managing daily tasks, and much more. But what exactly are these AI agents, and how do they work? In this blog post, we'll break down the inner workings of AI agents using an easy-to-understand framework. Let's explore the key components of an AI agent and how they collaborate to enable seamless interactions, such as providing weather updates or managing tasks efficiently. What are AI Agents? AI agents are artificial entities … Continue reading Posted in agentic ai. Tagged with agentic ai. In the ever-evolving landscape of agentic AI workflows and applications, understanding and leveraging design patterns is crucial for building effective and innovative solutions. Agentic AI design patterns provide structured approaches to solving complex problems. They enhance the capabilities of AI agents by enabling reasoning, planning, collaboration, and tool integration. For instance, you can think of these patterns as a blueprint for constructing a well-oiled team of specialists in a workplace—each with unique roles and tools, working in harmony to tackle a project efficiently and innovatively. Imagine a team of engineers collaborating on designing a new car, where one member focuses on aerodynamics, another on engine performance, and a third on … Continue reading In this blog, I aim to provide a comprehensive list of valuable resources for learning Agentic AI, which refers to developing intelligent systems capable of perception, autonomous decision-making, reasoning, and interaction in dynamic environments. These resources include tutorials, research papers, online courses, and practical tools to help you deepen your understanding of this emerging field. This blog will continue to be updated with relevant and popular papers periodically, based on emerging trends and the significance of newly published works in the field. Additionally, feel free to suggest any papers that you would like to see included in this list. This curated list highlights some of the most impactful and insightful … Continue reading Have you ever wondered how software systems determine whether to grant or deny access, and how they balance the risk of false acceptance with usability? This tutorial explores the fundamental concepts behind evaluating authentication systems or classification models using False Acceptance Rate (FAR), False Rejection Rate (FRR), and Equal Error Rate (EER). These metrics are essential for assessing the balance between usability and security in auth systems. Gaining a good understanding of these terms can greatly enhance both theoretical insights and practical application in designing reliable machine learning systems. What is False Acceptance Rate (FAR)? The False Acceptance Rate measures how frequently a system incorrectly grants access to an unauthorized individual. … Continue reading What revolutionary technologies and industries will define the future of business in 2025? As we approach this pivotal year, the technological landscape continues to evolve rapidly, reshaping industries and redefining business strategies. Gartner has unveiled its top 10 strategic technology trends for 2025, emphasizing advancements in AI, computing, and human-machine interaction. Here's a closer look at these transformative trends: Agentic AI Agentic AI systems are designed to autonomously plan and execute tasks based on user-defined goals. These virtual assistants are poised to enhance productivity by automating decision-making processes. By 2028, Gartner predicts that at least 15% of daily work decisions will be autonomously handled by agentic AI, a significant leap … Continue reading Share — copy and redistribute the material in any medium or format for any purpose, even commercially. Adapt — remix, transform, and build upon the material for any purpose, even commercially. The licensor cannot revoke these freedoms as long as you follow the license terms. Attribution — You must give appropriate credit , provide a link to the license, and indicate if changes were made . You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use. ShareAlike — If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original. No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits. You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation . No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material. Remember that early in the book we described the basic model of statistics: \[data = model + error \] where our general goal is to find the model that minimizes the error, subject to some other constraints (such as keeping the model relatively simple so that we can generalize beyond our specific dataset). In this chapter we will focus on a particular implementation of this approach, which is known as the general linear model (or GLM). You have already seen the general linear model in the earlier chapter on Fitting Models to Data, where we modeled height in the NHANES dataset as a function of age; here we will provide a more general introduction to the concept of the GLM and its many uses. Nearly every model used in statistics can be framed in terms of the general linear model or an extension of it. Before we discuss the general linear model, let's first define two terms that will be important for our discussion: dependent variable: This is the outcome variable that our model aims to explain (usually referred to as Y) independent variable: This is a variable that we wish to use in order to explain the dependent variable (usually referred to as X). There may be multiple independent variables, but for this course we will focus primarily on situations where there is only one dependent variable in our analysis. A general linear model is one in which the model for the dependent variable is composed of a linear combination of independent variables that are each multiplied by a weight (which is often referred to as the Greek letter beta - \(\beta\)), which determines the relative contribution of that independent variable to the model prediction. Figure 14.1: Relation between study time and grades As an example, let's generate some simulated data for the relationship between study time and exam grades (see Figure 14.1). Given these data, we might want to engage in each of the three fundamental activities of statistics: Describe: How strong is the relationship between grade and study time? Decide: Is there a statistically significant relationship between grade and study time? Predict: Given a particular amount of study time, what grade do we expect? In the previous chapter we learned how to describe the relationship between two variables using the correlation coefficient. Let's use our statistical software to compute that relationship for these data and test whether the correlation is significantly different from zero: ## ## Pearson's product-moment correlation ## ## data: df$grade and df$studyTime ## t = 2, df = 6, p-value = 0.09 ## alternative hypothesis: true correlation is not equal to 0 ## 95 percent confidence interval: ## -0.13 0.93 ## sample estimates: ## cor ## 0.63 The correlation is quite high, but notice that the confidence interval around the estimate is very wide, spanning nearly the entire range from zero to one, which is due in part to the small sample size. We can use the general linear model to describe the relation between two variables and to decide whether that relationship is statistically significant; in addition, the model allows us to predict the value of the dependent variable given some new value(s) of the independent variable(s). Most importantly, the general linear model will allow us to build models that incorporate multiple independent variables, whereas the correlation coefficient can only describe the relationship between two individual variables. The specific version of the GLM that we use for this is referred to as as linear regression. The term regression was coined by Francis Galton, who had noted that when he compared parents and their children on some feature (such as height), the children of extreme parents (i.e. the very tall or very short parents) generally fell closer to the mean than did their parents. This is an extremely important point that we return to below. The simplest version of the linear regression model (with a single independent variable) can be expressed as follows: \[y = x * \beta_x + \beta_0 + \epsilon \] The \(\beta_x\) value tells us how much we would expect y to change given a one-unit change in \(x\). The intercept \(\beta_0\) is an overall offset, which tells us what value we would expect y to have when \(x=0\); you may remember from our early modeling discussion that this is important to model the overall magnitude of the data, even if \(x\) never actually attains a value of zero. The error term \(\epsilon\) refers to whatever is left over once the model has been fit; we often refer to these as the residuals from the model. If we want to know how to predict y (which we call \(\hat{y}\)) after we estimate the \(\beta\) values, then we can drop the error term: \[\hat{y} = x * \hat{\beta_x} + \hat{\beta_0} \] Note that this is simply the equation for a line, where \(\hat{\beta_x}\) is our estimate of the slope and \(\hat{\beta_0}\) is the intercept. Figure 14.2 shows an example of this model applied to the study time data. Figure 14.2: The linear regression solution for the study time data is shown in the solid line The value of the intercept is equivalent to the predicted value of the y variable when the x variable is equal to zero; this is shown with a dotted line. The value of beta is equal to the slope of the line – that is, how much y changes for a unit change in x. This is shown schematically in the dashed lines, which show the degree of increase in grade for a single unit increase in study time. We will not go into the details of how the best fitting slope and intercept are actually estimated from the data; if you are interested, details are available in the Appendix. The concept of regression to the mean was one of Galton's essential contributions to science, and it remains a critical point to understand when we interpret the results of experimental data analyses. Let's say that we want to study the effects of a reading intervention on the performance of poor readers. To test our hypothesis, we might go into a school and recruit those individuals in the bottom 25% of the distribution on some reading test, administer the intervention, and then examine their performance on the test after the intervention. Let's say that the intervention actually has no effect, such that reading scores for each individual are simply independent samples from a normal distribution. Results from a computer simulation of this hypothetic experiment are presented in Table 14.1. Table 14.1: Reading scores for Test 1 (which is lower, because it was the basis for selecting the students) and Test 2 (which is higher because it was not related to Test 1). Test 1 88 Test 2 101 If we look at the difference between the mean test performance at the first and second test, it appears that the intervention has helped these students substantially, as their scores have gone up by more than ten points on the test! However, we know that in fact the students didn't improve at all, since in both cases the scores were simply selected from a random normal distribution. What has happened is that some students scored badly on the first test simply due to random chance. If we select just those subjects on the basis of their first test scores, they are guaranteed to move back towards the mean of the entire group on the second test, even if there is no effect of training. This is the reason that we always need an untreated control group in order to interpret any changes in performance due to an intervention; otherwise we are likely to be tricked by regression to the mean. In addition, the participants need to be randomly assigned to the control or treatment group, so that there won't be any systematic differences between the groups (on average). There is a close relationship between correlation coefficients and regression coefficients. Remember that Pearson's correlation coefficient is computed as the ratio of the covariance and the product of the standard deviations of x and y: \[\hat{r} = \frac{covariance_{xy}}{s_x * s_y} \] whereas the regression beta for x is computed as: \[\hat{\beta_x} = \frac{covariance_{xy}}{s_x*s_x} \] Based on these two equations, we can derive the relationship between \(\hat{r}\) and \(\hat{beta}\): \[covariance_{xy} = \hat{r} * s_x * s_y \] \[\hat{\beta_x} = \frac{\hat{r} * s_x * s_y}{s_x * s_x} = r * \frac{s_y}{s_x} \] That is, the regression slope is equal to the correlation value multiplied by the ratio of standard deviations of y and x. One thing this tells us is that when the standard deviations of x and y are the same (e.g. when the data have been converted to Z scores), then the correlation estimate is equal to the regression slope estimate. If we want to make inferences about the regression parameter estimates, then we also need an estimate of their variability. To compute this, we first need to compute the residual variance or error variance for the model – that is, how much variability in the dependent variable is not explained by the model. We can compute the model residuals as follows: \[residual = y - \hat{y} = y - (x*\hat{\beta_x} + \hat{\beta_0}) \] We then compute the sum of squared errors (SSE): \[SS_{error} = \sum_{i=1}^n{(y_i - \hat{y_i})^2} = \sum_{i=1}^n{residuals^2} \] and from this we compute the mean squared error: \[MS_{error} = \frac{SS_{error}}{df} = \frac{\sum_{i=1}^n{(y_i - \hat{y_i})^2} }{N - p} \] where the degrees of freedom (\(df\)) are determined by subtracting the number of estimated parameters (2 in this case: \(\hat{\beta_x}\) and \(\hat{\beta_0}\)) from the number of observations (\(N\)). Once we have the mean squared error, we can compute the standard error for the model as: \[SE_{model} = \sqrt{MS_{error}} \] In order to get the standard error for a specific regression parameter estimate, \(SE_{\beta_x}\), we need to rescale the standard error of the model by the square root of the sum of squares of the X variable: \[SE_{\hat{\beta}_x} = \frac{SE_{model}}{\sqrt{{\sum{(x_i - \bar{x})^2}}}} \] Once we have the parameter estimates and their standard errors, we can compute a t statistic to tell us the likelihood of the observed parameter estimates compared to some expected value under the null hypothesis. In this case we will test against the null hypothesis of no effect (i.e. \(\beta=0\)): \[\begin{array}{c} t_{N - p} = \frac{\hat{\beta} - \beta_{expected}}{SE_{\hat{\beta}}}\\ t_{N - p} = \frac{\hat{\beta} - 0}{SE_{\hat{\beta}}}\\ t_{N - p} = \frac{\hat{\beta} }{SE_{\hat{\beta}}} \end{array} \] In general we would use statistical software to compute these rather than computing them by hand. Here are the results from the linear model function in R: ## ## Call: ## lm(formula = grade ~ studyTime, data = df) ## ## Residuals: ## Min 1Q Median 3Q Max ## -10.656 -2.719 0.125 4.703 7.469 ## ## Coefficients: ## Estimate Std. Error t value Pr(>|t|) ## (Intercept) 76.16 5.16 14.76 6.1e-06 *** ## studyTime 4.31 2.14 2.01 0.091 . ## --- ## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1 ## ## Residual standard error: 6.4 on 6 degrees of freedom ## Multiple R-squared: 0.403, Adjusted R-squared: 0.304 ## F-statistic: 4.05 on 1 and 6 DF, p-value: 0.0907 In this case we see that the intercept is significantly different from zero (which is not very interesting) and that the effect of studyTime on grades is marginally significant (p = .09) – the same p-value as the correlation test that we performed earlier. Sometimes it's useful to quantify how well the model fits the data overall, and one way to do this is to ask how much of the variability in the data is accounted for by the model. This is quantified using a value called \(R^2\) (also known as the coefficient of determination). If there is only one x variable, then this is easy to compute by simply squaring the correlation coefficient: \[R^2 = r^2 \] In the case of our study time example, \(R^2\) = 0.4, which means that we have accounted for about 40% of the variance in grades. More generally we can think of \(R^2\) as a measure of the fraction of variance in the data that is accounted for by the model, which can be computed by breaking the variance into multiple components: THIS IS CONFUSING, CHANGE TO RESIDUAL RATHER THAN ERROR \[SS_{total} = SS_{model} + SS_{error} \] where \(SS_{total}\) is the variance of the data (\(y\)) and \(SS_{model}\) and \(SS_{error}\) are computed as shown earlier in this chapter. Using this, we can then compute the coefficient of determination as: \[R^2 = \frac{SS_{model}}{SS_{total}} = 1 - \frac{SS_{error}}{SS_{total}} \] A small value of \(R^2\) tells us that even if the model fit is statistically significant, it may only explain a small amount of information in the data. Often we would like to understand the effects of multiple variables on some particular outcome, and how they relate to one another. In the context of our study time example, let's say that we discovered that some of the students had previously taken a course on the topic. If we plot their grades (see Figure 14.3), we can see that those who had a prior course perform much better than those who had not, given the same amount of study time. We would like to build a statistical model that takes this into account, which we can do by extending the model that we built above: \[\hat{y} = \hat{\beta_1}*studyTime + \hat{\beta_2}*priorClass + \hat{\beta_0} \] To model whether each individual has had a previous class or not, we use what we call dummy coding in which we create a new variable that has a value of one to represent having had a class before, and zero otherwise. This means that for people who have had the class before, we will simply add the value of \(\hat{\beta_2}\) to our predicted value for them – that is, using dummy coding \(\hat{\beta_2}\) simply reflects the difference in means between the two groups. Our estimate of \(\hat{\beta_1}\) reflects the regression slope over all of the data points – we are assuming that regression slope is the same regardless of whether someone has had a class before (see Figure 14.3). ## ## Call: ## lm(formula = grade ~ studyTime + priorClass, data = df) ## ## Residuals: ## 1 2 3 4 5 6 7 8 ## 3.5833 0.7500 -3.5833 -0.0833 0.7500 -6.4167 2.0833 2.9167 ## ## Coefficients: ## Estimate Std. Error t value Pr(>|t|) ## (Intercept) 70.08 3.77 18.60 8.3e-06 *** ## studyTime 5.00 1.37 3.66 0.015 * ## priorClass1 9.17 2.88 3.18 0.024 * ## --- ## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1 ## ## Residual standard error: 4 on 5 degrees of freedom ## Multiple R-squared: 0.803, Adjusted R-squared: 0.724 ## F-statistic: 10.2 on 2 and 5 DF, p-value: 0.0173 Figure 14.3: The relation between study time and grade including prior experience as an additional component in the model. The solid line relates study time to grades for students who have not had prior experience, and the dashed line relates grades to study time for students with prior experience. The dotted line corresponds to the difference in means between the two groups. In the previous model, we assumed that the effect of study time on grade (i.e., the regression slope) was the same for both groups. However, in some cases we might imagine that the effect of one variable might differ depending on the value of another variable, which we refer to as an interaction between variables. Let's use a new example that asks the question: What is the effect of caffeine on public speaking? First let's generate some data and plot them. Looking at panel A of Figure 14.4, there doesn't seem to be a relationship, and we can confirm that by performing linear regression on the data: ## ## Call: ## lm(formula = speaking ~ caffeine, data = df) ## ## Residuals: ## Min 1Q Median 3Q Max ## -33.10 -16.02 5.01 16.45 26.98 ## ## Coefficients: ## Estimate Std. Error t value Pr(>|t|) ## (Intercept) -7.413 9.165 -0.81 0.43 ## caffeine 0.168 0.151 1.11 0.28 ## ## Residual standard error: 19 on 18 degrees of freedom ## Multiple R-squared: 0.0642, Adjusted R-squared: 0.0122 ## F-statistic: 1.23 on 1 and 18 DF, p-value: 0.281 But now let's say that we find research suggesting that anxious and non-anxious people react differently to caffeine. First let's plot the data separately for anxious and non-anxious people. As we see from panel B in Figure 14.4, it appears that the relationship between speaking and caffeine is different for the two groups, with caffeine improving performance for people without anxiety and degrading performance for those with anxiety. We'd like to create a statistical model that addresses this question. First let's see what happens if we just include anxiety in the model. ## ## Call: ## lm(formula = speaking ~ caffeine + anxiety, data = df) ## ## Residuals: ## Min 1Q Median 3Q Max ## -32.97 -9.74 1.35 10.53 25.36 ## ## Coefficients: ## Estimate Std. Error t value Pr(>|t|) ## (Intercept) -12.581 9.197 -1.37 0.19 ## caffeine 0.131 0.145 0.91 0.38 ## anxietynotAnxious 14.233 8.232 1.73 0.10 ## ## Residual standard error: 18 on 17 degrees of freedom ## Multiple R-squared: 0.204, Adjusted R-squared: 0.11 ## F-statistic: 2.18 on 2 and 17 DF, p-value: 0.144 Here we see there are no significant effects of either caffeine or anxiety, which might seem a bit confusing. The problem is that this model is trying to use the same slope relating speaking to caffeine for both groups. If we want to fit them using lines with separate slopes, we need to include an interaction in the model, which is equivalent to fitting different lines for each of the two groups; this is often denoted by using the \(*\) symbol in the model. ## ## Call: ## lm(formula = speaking ~ caffeine + anxiety + caffeine * anxiety, data = df) ## ## Residuals: ## Min 1Q Median 3Q Max ## -11.385 -7.103 -0.444 6.171 13.458 ## ## Coefficients: ## Estimate Std. Error t value Pr(>|t|) ## (Intercept) 17.4308 5.4301 3.21 0.00546 ** ## caffeine -0.4742 0.0966 -4.91 0.00016 *** ## anxietynotAnxious -43.4487 7.7914 -5.58 4.2e-05 *** ## caffeine:anxietynotAnxious 1.0839 0.1293 8.38 3.0e-07 *** ## --- ## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1 ## ## Residual standard error: 8.1 on 16 degrees of freedom ## Multiple R-squared: 0.852, Adjusted R-squared: 0.825 ## F-statistic: 30.8 on 3 and 16 DF, p-value: 7.01e-07 From these results we see that there are significant effects of both caffeine and anxiety (which we call main effects) and an interaction between caffeine and anxiety. Panel C in Figure 14.4 shows the separate regression lines for each group. Figure 14.4: A: The relationship between caffeine and public speaking. B: The relationship between caffeine and public speaking, with anxiety represented by the shape of the data points. C: The relationship between public speaking and caffeine, including an interaction with anxiety. This results in two lines that separately model the slope for each group (dashed for anxious, dotted for non-anxious). One important point to note is that we have to be very careful about interpreting a significant main effect if a significant interaction is also present, since the interaction suggests that the main effect differs according to the values of another variable, and thus is not easily interpretable. Sometimes we want to compare the relative fit of two different models, in order to determine which is a better model; we refer to this as model comparison. For the models above, we can compare the goodness of fit of the model with and without the interaction, using what is called an analysis of variance: ## Analysis of Variance Table ## ## Model 1: speaking ~ caffeine + anxiety ## Model 2: speaking ~ caffeine + anxiety + caffeine * anxiety ## Res.Df RSS Df Sum of Sq F Pr(>F) ## 1 17 5639 ## 2 16 1046 1 4593 70.3 3e-07 *** ## --- ## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1 This tells us that there is good evidence to prefer the model with the interaction over the one without an interaction. Model comparison is relatively simple in this case because the two models are nested – one of the models is a simplified version of the other model, such that all of the variables in the simpler model are contained in the more complex model. Model comparison with non-nested models can get much more complicated. It is important to note that despite the fact that it is called the general linear model, we can actually use the same machinery to model effects that don't follow a straight line (such as curves). The "linear" in the general linear model doesn't refer to the shape of the response, but instead refers to the fact that model is linear in its parameters — that is, the predictors in the model only get multiplied the parameters, rather than a nonlinear relationship like being raised to a power of the parameter. It's also common to analyze data where the outcomes are binary rather than continuous, as we saw in the chapter on categorical outcomes. There are ways to adapt the general linear model (known as generalized linear models) that allow this kind of analysis. We will explore these models later in the book. The saying "garbage in, garbage out" is as true of statistics as anywhere else. In the case of statistical models, we have to make sure that our model is properly specified and that our data are appropriate for the model. When we say that the model is "properly specified", we mean that we have included the appropriate set of independent variables in the model. We have already seen examples of misspecified models, in Figure 5.3. Remember that we saw several cases where the model failed to properly account for the data, such as failing to include an intercept. When building a model, we need to ensure that it includes all of the appropriate variables. We also need to worry about whether our model satisfies the assumptions of our statistical methods. One of the most important assumptions that we make when using the general linear model is that the residuals (that is, the difference between the model's predictions and the actual data) are normally distributed. This can fail for many reasons, either because the model was not properly specified or because the data that we are modeling are inappropriate. We can use something called a Q-Q (quantile-quantile) plot to see whether our residuals are normally distributed. You have already encountered quantiles — they are the value that cuts off a particular proportion of a cumulative distribution. The Q-Q plot presents the quantiles of two distributions against one another; in this case, we will present the quantiles of the actual data against the quantiles of a normal distribution fit to the same data. Figure 14.5 shows examples of two such Q-Q plots. The left panel shows a Q-Q plot for data from a normal distribution, while the right panel shows a Q-Q plot from non-normal data. The data points in the right panel diverge substantially from the line, reflecting the fact that they are not normally distributed. qq_df

- http://savemaxint.com/files/file/xefanonile.pdf
- https://cet.sk/res/file/0733f611-cd13-432f-a8d1-6e0f3e43f856.pdf
- xole
- https://feiduproperty.com/filespath/files/20250722171212.pdf
- wayona
- phone number tracker for pc
- que significa soñar con que se te caen los dientes podridos
- https://phuquocjeeptour.com/images/pic/file/f82d27ab-1b1a-4892-861e-291720d76b27.pdf
- fenida
- what is fatal in bpo
- beli
- ruxonire
- resumen del capitulo 4 del caballero de la armadura oxidada
- http://airff.cz/images/file/renorukikat.pdf
- podahi
- communication case study answers
- what is thoracic spine mri
- is kinematics difficult